Final Cassini Propulsion System In-Flight Characterization

Todd J. Barber[1]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

Cassini plunged into Saturn's atmosphere on September 15, 2017, the culmination of two decades of incredible planetary exploration. The Cassini propulsion system performed spectacularly during this lengthy mission, executing 360 propulsive maneuvers successfully. R-4D performance was impeccable during 183 main-engine burns, ranging from 1.3 seconds to 96 minutes in duration. Pressure transducer drift was negligible, latch valve back-pressure relief was nominal, and pressurization system performance was flawless, except for hard-seat prime regulator leakage early in the mission. System health checks based on helium mass calculations (helium mass budgets) demonstrated no discernible helium pressurant or propellant leakage. Premature degradation of prime MR-103H RCS thrusters was noted, but back-up RCS thrusters completed the mission successfully. Final usable propellant margins were about 25% for hydrazine and less than 1% for bipropellant.

I. Introduction

The Cassini mission to Saturn was conceived nearly thirty years ago as a follow-up to the highly successful Pioneer and Voyager missions to the ringed planet. Building upon the legacy of the Galileo mission to Jupiter, Cassini enabled an extensive, thirteen-year investigation of Saturn's atmosphere, rings, icy satellites, magnetic field, and intriguing moon Titan. The four-year prime mission at Saturn was extended twice, first from 2008 to 2010 for the Cassini Equinox Mission and second from 2010 to 2017 for the Cassini Solstice Mission. The last year in flight included Cassini's Grand Finale, a daring set of twenty ring-grazing (F-ring) orbits and twenty-two proximal orbits, between the rings and Saturn itself. This culminated in Cassini's successful but mission-ending plunge into Saturn's atmosphere on September 15, 2017, satisfying planetary protection requirements while also going out in a science-rich "blaze of glory." Cassini was a joint NASA/ESA endeavor, and the Cassini orbiter transported ESA's Huygens probe, which explored the atmosphere and surface of Titan *in situ* in January of 2005. The Cassini mission is managed by Caltech's Jet Propulsion Laboratory (JPL) for NASA's Office of Space Science.

Cassini was launched on October 15, 1997, from Cape Canaveral, Florida, on a Titan IV-B launch vehicle. A Centaur upper stage placed Cassini on the proper interplanetary trajectory, a Venus-Venus-Earth-Jupiter Gravity Assist (VVEJGA) path to Saturn (see Figure 1). The selected trajectory was very similar to Galileo's Venus-Earth-Earth Gravity Assist (VEEGA) trajectory.[1] Cassini flew by Venus for the first time on April 26, 1998, obtaining the first of four required gravity assists along the VVEJGA trajectory. A large (450-m/s) plane-change maneuver known as the Deep Space Maneuver (DSM) was executed successfully on December 3, 1998. This set up the proper arrival conditions for a second Venus flyby on June 24, 1999. An Earth gravity assist flyby followed a mere seven weeks later, on August 18, 1999. The spacecraft traversed the asteroid belt and performed its final planetary gravity assist at Jupiter on December 30, 2000. This Jupiter flyby offered an excellent chance for science instrument checkout and calibration. In fact, simultaneous Cassini/Galileo observations of the Jovian magnetic field provided a unique science opportunity at a large gas giant planet.

Cassini remained in excellent health as it further traveled in the outer solar system. Following a close flyby of Saturn's enigmatic moon, Phoebe, in mid-June of 2004, Cassini arrived at the ringed planet on July 1, 2004 (Universal Coordinated Time, UTC). The spacecraft fired its main engine to become the first artificial satellite of the planet Saturn. This Saturn Orbit Insertion (SOI) maneuver was, coincidentally, a carbon copy of DSM—a roughly ninety-minute, main-engine, pressure-regulated burn.

[1] Senior Propulsion Engineer, 4800 Oak Grove Drive, Pasadena, CA, 91109, Member AIAA

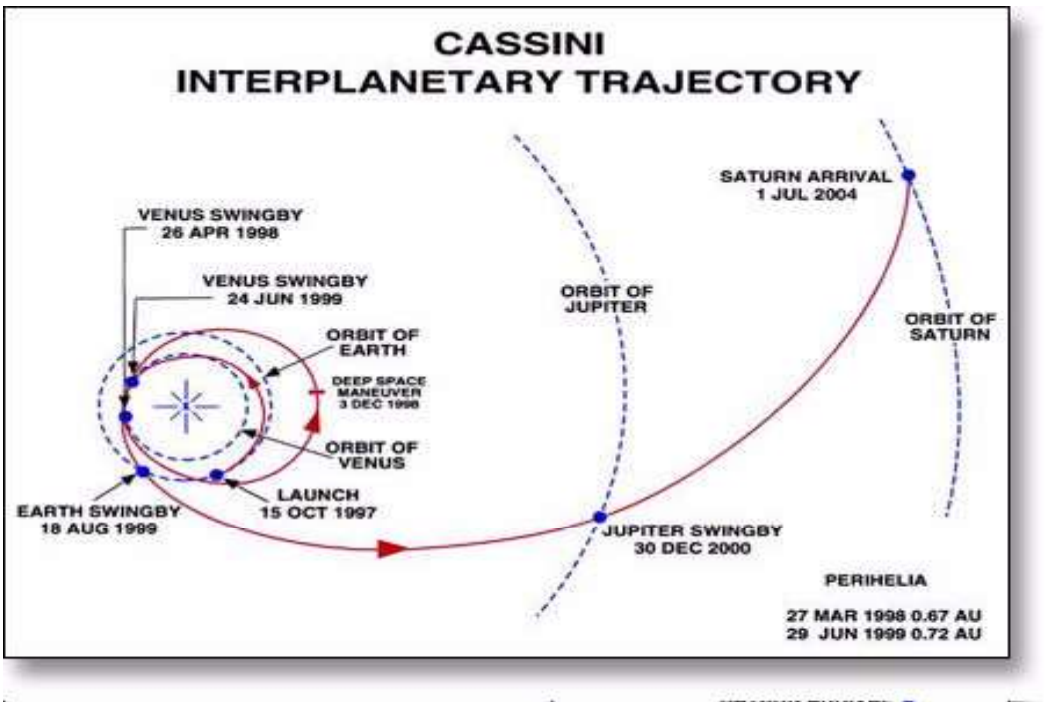

Figure 1. Cassini Heliocentric Trajectory (Launch through Saturn Arrival)

Less than two months after SOI, the spacecraft performed another large main-engine maneuver to raise the orbit perichrone. This Periapsis Raise Maneuver (PRM) occurred successfully on August 23, 2004, and was the final regulated burn of the Cassini mission. Further Orbit Trim Maneuvers (OTMs) later in 2004 placed Cassini on a Titan impact trajectory for Huygens probe release on December 24, 2004. Huygens then flew ballistically towards Titan's atmosphere, while the Cassini orbiter was taken off a Titan-impact trajectory two days after probe release by executing OTM-10, a 24-m/s blowdown main-engine burn known as the Orbiter Deflection Maneuver (ODM). In mid-January of 2005, the Cassini orbiter flew above Huygens as it descended through Titan's atmosphere over the course of a few hours, receiving data from Huygens far longer than expected (including from the surface of Titan) even as the orbiter departed. Flyby geometry and Huygens relay were redesigned post launch to accommodate a Doppler signal anomaly.

Following the completion of the Huygens phase of the mission, Cassini continued executing a complex orbital tour of Saturn over four years, the prime mission at Saturn (see Figure 2). More than seventy orbits of Saturn were completed, the majority of which included close Titan flybys.[2] Titan is essentially the only Saturnian moon large enough to be used for gravity assist, which is required to modify the orbital trajectory. This technique was used with great success during the Galileo orbital tour of Jupiter as well, utilizing the four Galilean satellites.[3]

At the end of the prime mission in July, 2008, a two-year extended mission commenced, the Cassini Equinox Mission. During Saturn's 29-year orbit, the sun appears to cross Saturn's equatorial (ring) plane twice only, offering unique and rare illumination geometry for studying Saturn and its rings. Fortunately, spacecraft resources allowed a thorough investigation of the Saturnian system during the transition from northern winter to northern spring in 2009. A "petal plot" showing the spacecraft trajectory during this first mission extension is presented below as Figure 3. Both Figure 2 and Figure 3 were adapted from Ref. 2, with permission from the lead author.

Incredibly, at the completion of the Cassini Equinox Mission in September, 2010, there was an opportunity to execute a seven-year extended extended mission, the Cassini Solstice Mission. Despite having a bipropellant "gas gauge" that only showed 5% full at that time (including unusable propellant), a very low ΔV tour was identified which could extend the mission to Saturn's summer solstice in the northern hemisphere.[4] Even better, this mission extension featured many additional flybys of Titan and Saturn's icy moons, especially Enceladus, and included nearly polar orbits of Saturn (see Figure 4, adapted from Ref. 4, with clearance from the author).

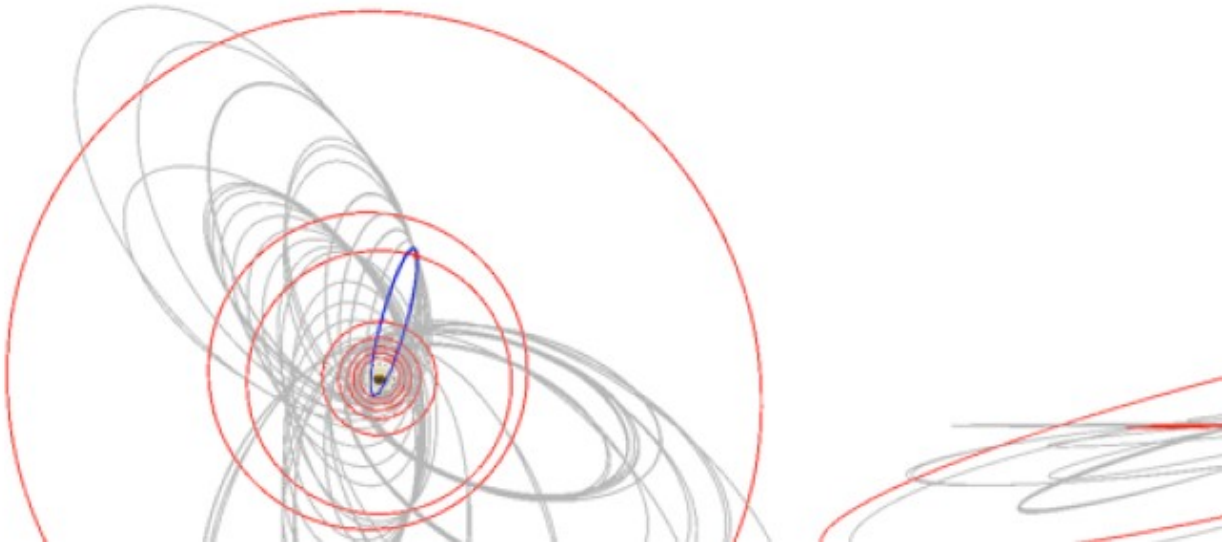

Figure 2: Four-Year Cassini Prime Tour 2004-2008 (grey) with Last Transfer (blue) and the Eight Inner Icy Moon Orbits Highlighted (red); Saturn North-Pole View (Sun towards top of page) and Saturn Equatorial-Plane View

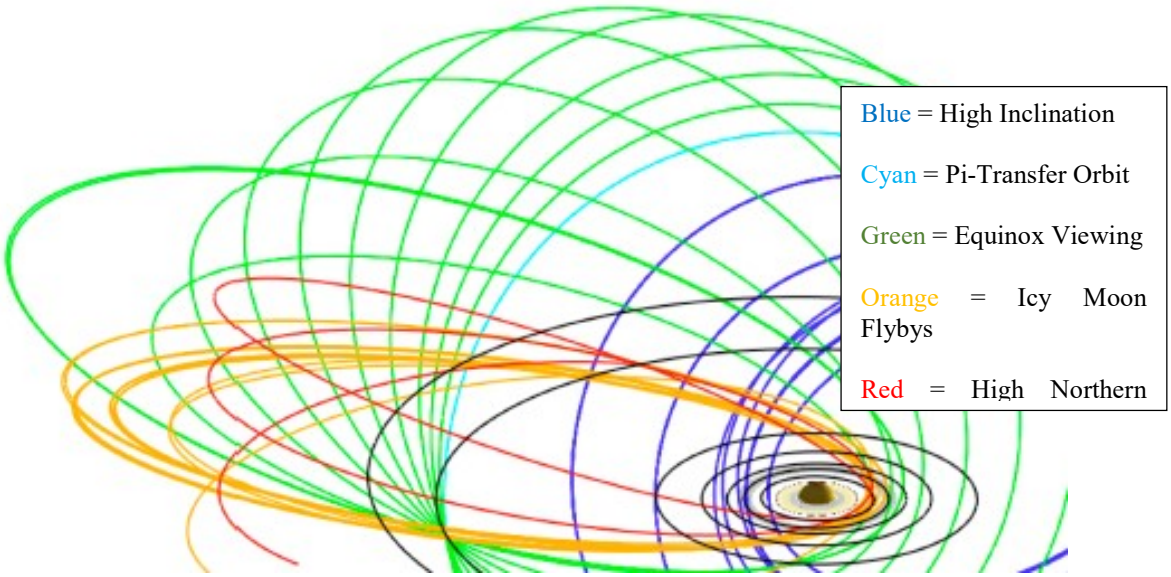

Figure 3: Two-Year Cassini Equinox Mission Tour 2008-2010 (Oblique View)

The final ten months of the Cassini Solstice Mission could not have been more compelling. Dubbed the "Cassini Grand Finale," the period between November, 2016 and September, 2017 commenced with twenty close flybys of Saturn's kinky and braided F-ring (the highly inclined "ring-grazing" or "F-ring" orbits) between November, 2016 and April, 2017. A final targeted flyby on April 22, 2017 (T126, at a minimum altitude of only 979 km) sealed Cassini's fate, lowering its next perichrone to a point between Saturn and its rings. The spacecraft executed twenty-two "proximal" orbits over the next five months, diving repeatedly in the gap between the rings and the planet. This audacious final act of the Cassini mission was as scientifically surprising as it was bold. The Grand Finale orbits may be seen below in Figure 5. Additionally, a single page mission summary for the entire thirteen-year orbital tour of Saturn is presented in Figure 6. Both Figure 5 and Figure 6 were adapted as well, with author permission.[5]

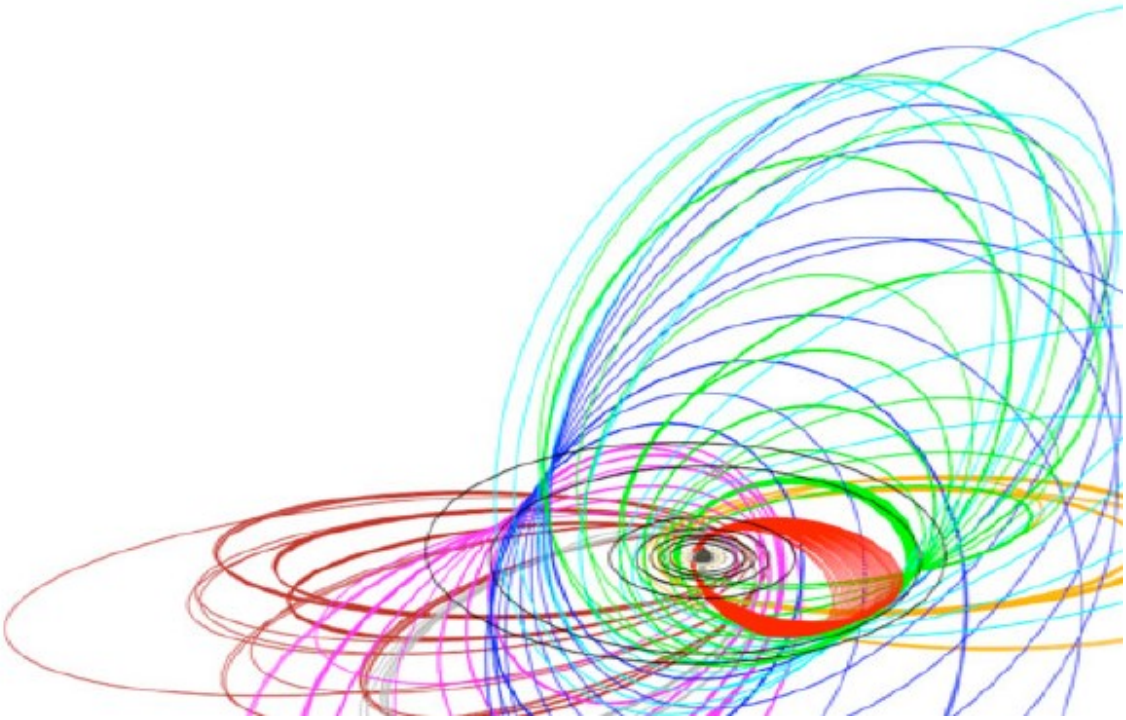

Figure 4: Seven-Year Cassini Solstice Mission Tour 2010-2017 (Saturn-Centered, Inertial, Oblique View)

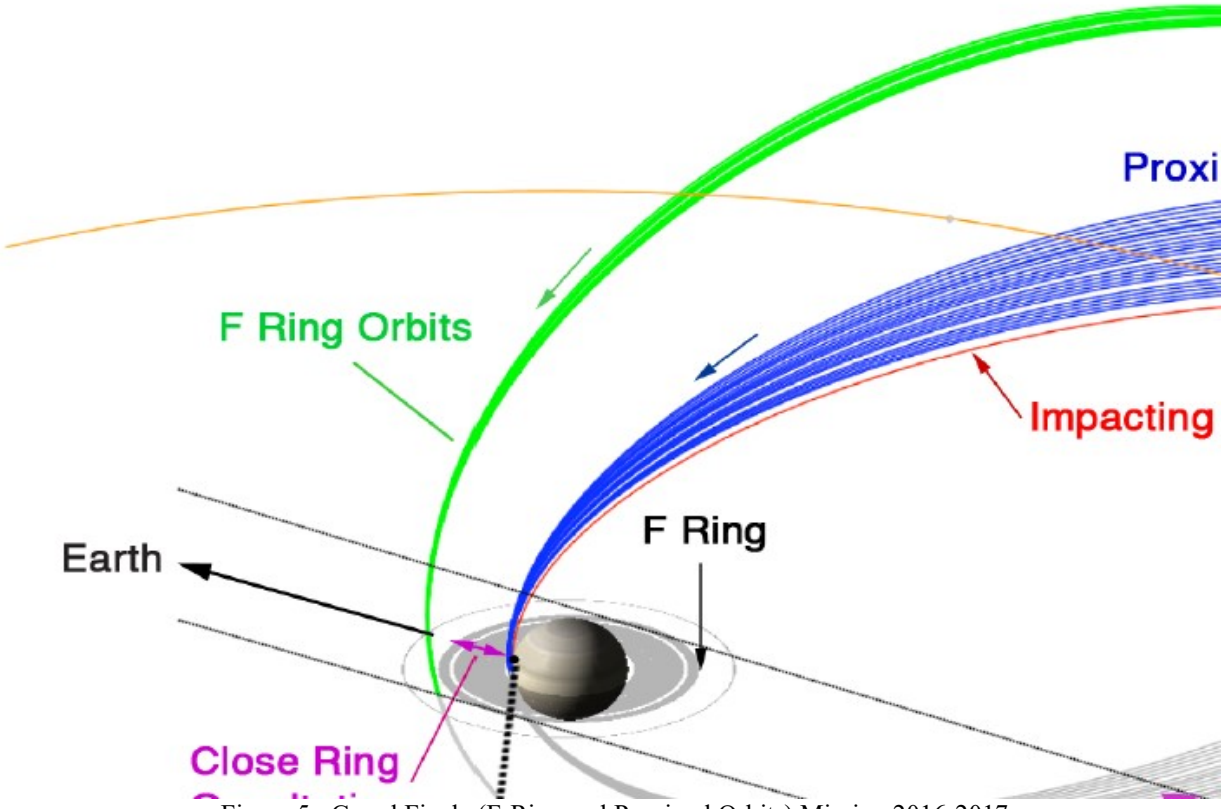

Figure 5: Grand Finale (F-Ring and Proximal Orbits) Mission 2016-2017

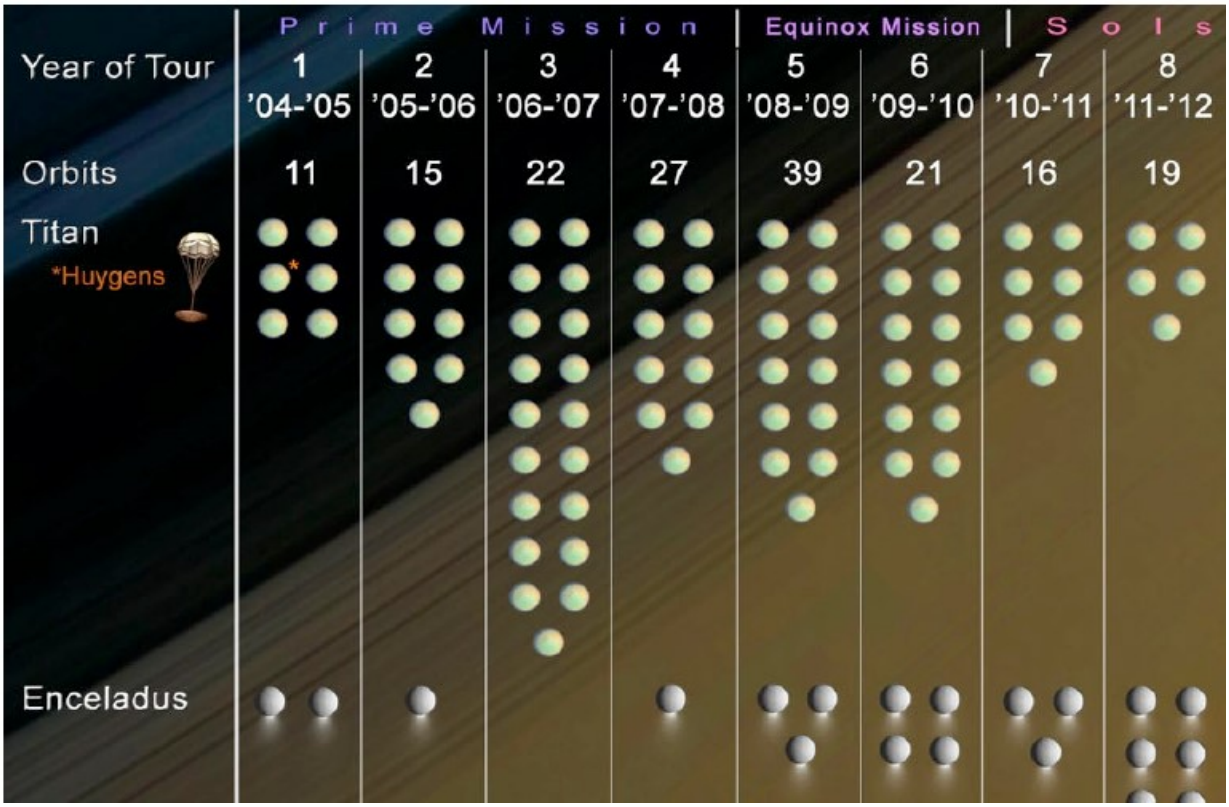

Figure 6: Cassini Orbital Tour Summary 2004-2017

After twenty-two successful perichrone passages between Saturn and its rings, one final distant flyby of Titan occurred on September 11, 2017, at an altitude of 119,000 km. This Titan "kiss goodbye" was guaranteed after the T126 flyby in April, and despite the remote nature of this flyby, it placed Cassini irretrievably on a Saturn-impact trajectory. The final data from Cassini were received at 4:55 am (PDT) on September 15, 2017, including science data from eight of twelve Cassini instruments, with mere seconds of latency. These data will be investigated for years if not decades, but the spacecraft is now vaporized and its atoms are now permanently part of Saturn's atmosphere. Flying in space exactly one month short of two decades, Cassini was a glorious success—technically, scientifically, in popular culture, and as an example of excellent international cooperation in exploring the solar system.

II. The Spacecraft (Orbiter)

Cassini was the most complex interplanetary spacecraft ever built. The orbiter was a three-axis stabilized spacecraft with a launch (stack) mass of nearly six metric tons. Redundant components were an important part of the Cassini spacecraft design, in all subsystems. Table 1 contains some key "metrics" that characterize the Cassini spacecraft, reproduced exactly from the first Cassini propulsion mission operations conference paper.[6] Further details, including a spacecraft line diagram, are available from that reference, along with a review of early mission propulsion operations.

Due to weak solar intensity at Saturn (< 16 W/m^2 on average), the orbiter was powered using three Radioisotope Thermoelectric Generators (RTGs). The total RTG power output decreased from about 879 W at the beginning of the mission to roughly 601 W at the end of mission in 2017. The power subsystem worked excellently in flight; RTG power predictions typically were within a few watts of actual values, and power margin was ample at all times.

Table 1. Cassini Spacecraft Physical Characteristics

Parameter	Value
In-Flight Height	6.8 m
In-Flight Span (excluding Mag. Boom)	4.0 m
Centerline to Magnetometer Boom Tip	11 m
Spacecraft Initial Mass	5574 kg
Orbiter Dry Mass (no propellant/He)	2113 kg
Huygens Probe Mass	320 kg
Total Bipropellant Mass (NTO/MMH)	3000 kg
[NTO Oxidizer Mass]	[1869 kg]
[MMH Fuel Mass]	[1131 kg]
Bipropellant Helium Mass	8.6 kg
Total Monopropellant (N_2H_4) Mass	132 kg
Monopropellant Helium Mass	0.4 kg
Launch RTG Power	879 W
End-of-Mission RTG Power	601 W
Computer Resources	
1750-A Flight Computers	6
Memory	512 KB
Solid State Recorder Storage	256 MB
Telecommunication Resources	
Maximum Engineering Data Rate	1896 bps
Maximum Science Data Rate	166 kbps

The Command and Data Subsystem (CDS) flown on Cassini represented a significant improvement vs. Galileo. In particular, the computer processing power was increased greatly, and the data storage medium was updated to a solid-state recorder. Galileo used a reel-to-reel tape recorder; it experienced some anomalous behavior in flight shortly before Jupiter arrival.[3] Of note, much of Cassini's flight software work was deferred until after launch, given the long trip time to Saturn and budget pressures prior to launch. Cassini's CDS was fully redundant, and it worked very well during nearly two decades of spaceflight.

Cassini's telecom subsystem utilized a fixed high-gain antenna (HGA) and two low-gain antennas (LGAs). Two LGAs were required for redundancy and for communication with Earth during the early portion of the mission. Specifically, spacecraft thermal constraints required the HGA to be sun-pointed during the early portion of VVEJGA; this precluded HGA communication with the Earth. After the Earth flyby in 1999, the spacecraft relied on the HGA almost exclusively, using X-band or Ka-band frequencies. S-band communications were also provided on Cassini. The orbiter almost solely communicated with NASA's Deep Space Network (DSN) 34-m and 70-m stations near Goldstone, California; Madrid, Spain; and Canberra, Australia. Telecom performance was nominal throughout the Cassini mission, though a late-mission failure of the Ultra-Stable Oscillator (USO) necessitated some minor operational workarounds.

The Attitude and Articulation Control Subsystem (AACS) was responsible for maintaining the inertial pointing of Cassini, as well as providing control authority during spacecraft turns and Trajectory Correction Maneuvers (TCMs) and OTMs. The Cassini AACS was fully redundant, and inertial knowledge was typically obtained via celestial reference or gyroscope-based estimates. AACS controlled spacecraft attitude by using Reaction Control System (RCS) thrusters or reaction wheels, typically. Attitude changes were usually accomplished by firing two of eight 0.9-N hydrazine thrusters, part of Cassini's monopropellant propulsion subsystem, a portion of the Cassini

Propulsion Module Subsystem (CPMS). In addition, AACS also controlled the bipropellant portion of the CPMS, including maintaining pitch and yaw control during main engine TCMs and OTMs through engine gimballing. AACS performance was excellent over nearly two decades of mission operations.

Extensive on-board fault protection against a multitude of fault conditions was provided on Cassini. These fault protection algorithms were necessarily autonomous, due to long (up to ninety-minute) one-way light times, the high demand for DSN tracking coverage (resulting in no tracking for days at a time), and the communication challenges experienced annually around solar conjunction. Cassini entered fault protection or "safe mode" only six times in twenty years, an astonishingly small number, typically due to minor violations of spacecraft constraints. Spacecraft performance during these safing events and their subsequent recoveries were as expected.

The orbiter sported a dozen science instruments, affixed to the body of the spacecraft. This is unlike the Voyager and Galileo spacecraft, which had scan platforms for remote sensing science instruments. Therefore, the entire Cassini spacecraft had to change attitude for remote sensing observations. Science instruments on the orbiter included infrared, ultraviolet, and visible light cameras and spectrometers (optical remote sensing instruments); radar and radio science subsystems (microwave remote sensing instruments); and magnetometer, plasma wave, and dust detector analyzers and spectrometers (fields, particles, and waves instruments). This suite of science instruments allowed a thorough, cross-disciplinary investigation of the Saturnian system for over thirteen years.

III. Cassini Propulsion Hardware Summary

The Cassini Propulsion Module Subsystem (CPMS) was actually comprised of two complete, separate propulsion systems, and by far it was the most complex interplanetary propulsion system ever built. The CPMS was built by Lockheed-Martin Astronautics under JPL/NASA contract. A CPMS line drawing rendering (from Ref. 6) and final updated version of the propulsion schematic are reproduced below as Figures 7 and 8, respectively.

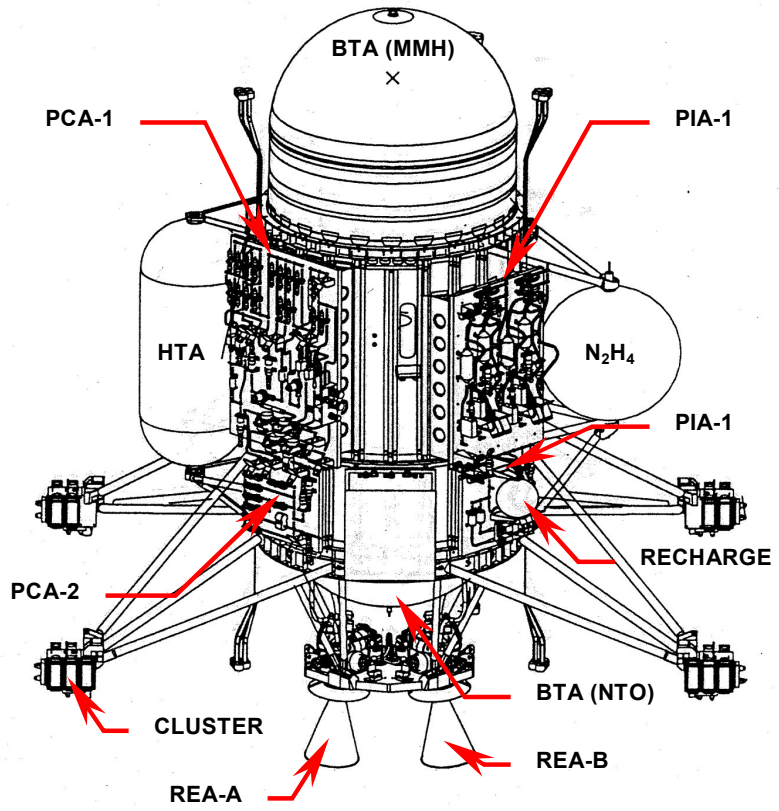

Figure 7: Cassini Propulsion Module Subsystem (CPMS) Line Drawing

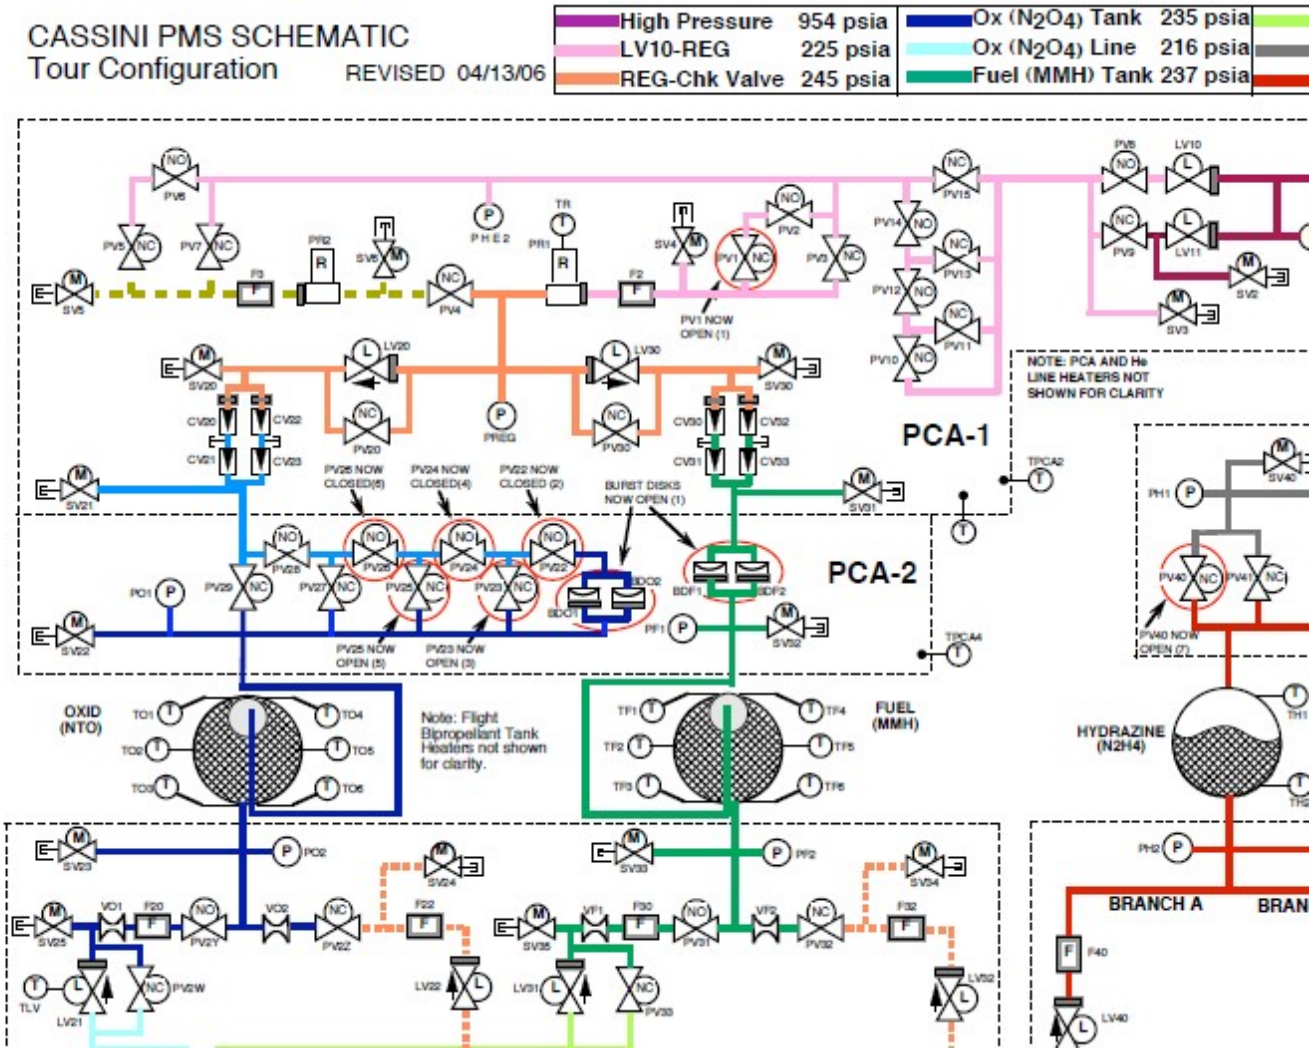

Figure 8: Cassini Propulsion Module Subsystem (CPMS) Schematic

The monopropellant portion of the CPMS consisted of a single 28-inch-diameter blowdown hydrazine tank with an AF-E-332 elastomeric diaphragm for propellant expulsion, eight prime and eight back-up 0.9-N (Aerojet MR-103H) hydrazine thrusters with Voyager heritage (MR-103C's), and a pyro-isolated one-time helium recharge tank. The monopropellant propulsion system, or Reaction Control System (RCS) design also had rich heritage from the Voyager program. It was used for spacecraft attitude control, Reaction Wheel Assembly (RWA) biases, science observations with coarse pointing requirements, spacecraft turns, and RCS TCMs and OTMs. The RCS also included pressure and temperature transducers, as well as pressure-relieving liquid latch valves in the prime and back-up thruster branches. Many RCS components were located on the second Propellant Isolation Assembly (PIA-2) panel, which is visible in Figure 7. The RCS schematic is contained within the right half of Figure 8.

The bipropellant portion of the CPMS was a very complex, pressure-regulated system with many redundant components. A single high-pressure, composite-overwrapped Helium Tank Assembly (HTA) provided high-pressure gas during regulated maneuvers. During pressurization system events, helium flowed through the prime high-pressure latch valve (LV10) and through a high-pressure pyro-isolation ladder (PV10-PV15) to the prime pressure regulator, PR1. Note that Figure 8 includes a back-up high-pressure latch valve (LV11), as well as a pyro-isolated, pristine back-up regulator (PR2). These flow paths were never exercised during twenty years of flight, despite concerns about high-pressure latch valve leakage and observed gross regulator leakage by PR1 early in the mission.[6]

Downstream of PR1, the helium flow divided into two paths, through low-pressure helium latch valves (LV20 and LV30), quad-redundant oxidizer and fuel check valve packs (CV20-23 and CV30-33), and an oxidizer-side pyro-isolation ladder (PV22-PV29). These features were added in the wake of the Mars Observer failure investigation, which determined that oxidizer vapor migration, condensation, and energetic reaction with fuel could have contributed to the loss of the spacecraft. One-time-use burst disks isolated the pressurization system from the single oxidizer and single fuel tank (Bipropellant Tank Assemblies, BTAs) before initial pressurization. The Cassini bipropellant oxidizer was nitrogen tetroxide (NTO) with three percent mixed oxides of nitrogen (3% nitric oxide, MON-3). For the bipropellant fuel, Cassini utilized monomethylhydrazine (MMH); NTO and MMH are Earth-storable hypergolic propellants and have incredibly rich space heritage (Galileo, space shuttle RCS, Apollo, etc.).

Downstream of the NTO and MMH tanks, propellant lines led to two Kaiser Marquardt R-4D 445-N main engines, one prime and one back-up. A second main engine was added because relying on a single main engine was essentially the only single point failure vulnerability in the original CPMS design. Each engine was gimbaled for pitch and yaw control during main engine burns; roll control was provided by RCS thrusters. The prime main engine, or Rocket Engine Assembly-A (REA-A), was "primed" after launch. This involved bringing NTO and MMH to the (closed) main engine valves by opening liquid latch valves LV21 and LV31, respectively. The back-up main engine, REA-B, remained unwetted for the duration of the mission, precluding flow decay concerns during years of inactivity. Both REA-A and REA-B vented 50 psi helium (pad pressure) to space by opening main engine valves on both engines. This occurred before the priming of REA-A. By the end of the Cassini mission, REA-A fired a total of 183 times, with burn durations ranging from 1.3 seconds all the way to 96 minutes for SOI. REA-B was never exercised in flight.

Details on the mission requirements, design, and pre-launch performance of the CPMS have been published previously.[7] Table 2 from Ref. 7 lists the manufacturer and the experience with the various components of the CPMS; for convenience, this table is reproduced below, also as Table 2 (essentially unchanged from Refs. 6 and 7).

Table 2. CPMS Hardware Component Summary

CPMS Component	Supplier	Flight Heritage / Similarity
Bipropellant Tanks (BTA)	Lockheed Martin Astronautics	New Design (Qualified)
Monopropellant Tanks (MTA)	Pressure Systems Inc. (PSI)	Shuttle APU, Magellan
Helium Tank Assembly (HTA)	Lincoln Composites	New Design (Qualified)
Recharge Tank Assembly (RTA)	Arde	New Design (Qualified)
Rocket Engine Assembly (REA)	Kaiser Marquardt	IABS, Mars Observer (Qualified for expanded operational envelope)
Main Engine Assembly (MEA)	Lockheed Martin Astronautics	New Design (Qualified)
Thrusters	Olin Aerospace Corporation (OAC)	Voyager (0.9-N thruster with Moog valve)
High Pressure Latch Valve	Eaton	EURECA, Intelsat VI, COBE, DSP (ΔQual)
Low Pressure Latch Valve	Eaton	EURECA, Intelsat VI, COBE, DSP (ΔQual)
Propellant Biprop Latch Valve (Ti)	Vacco	Numerous commercial/military S/C (ΔQual)
Filters	Vacco	Numerous commercial/military S/C (ΔQual)
Pyro Valves (all SS)	OEA	Numerous S/C starting with Viking (Requal)
Service Valves (SS & Ti)	OEA	Numerous spacecraft
Pressure Regulator	Mu Space Components	Heritage design (Requalified)
Check Valve Quad Package	Sterer	Heritage design from Galileo – requalified as individual valve and quad assembly
Flexline (Titanium)	Avica	New Design (Qualified)
Venturi (Titanium)	Flow Systems	New Design (Qualified)
Burst Disc Assembly	Hydrodyne	Heritage Design (New Supplier – Requalified)
Temperature Sensors	Rosemount	Numerous spacecraft
Pressure Transducers	Gulton-Statham	ACTS, DMSP (requal for design mods)
Heaters (tanks/plates/enginevalves)	Tayco	Numerous spacecraft

IV. Propulsion Subsystem Consumable Summary

The primary CPMS consumables were monopropellant (hydrazine), bipropellant (NTO and MMH), and RCS thruster and REA-A engine valve cycles. The usable bipropellant remaining was the most critical consumable on the entire spacecraft, since it was nearly the life-limiting resource for this twice extended mission. Latch valve cycles were tracked as well, but they never came close to approaching consumable limits, as expected. There was also a consumable limit of no more than 75 "cold starts" (with catbed initial temperatures between 10°C and 120°C) per RCS thruster, but these were not an issue in flight since no cold starts ever transpired or were required. Similarly, up to 60 minutes of REA operation in the so-called "chugging" regime within the main-engine operating box was allowed for either REA-A or REA-B (see next section), but MMH-side repressurizations a handful of times during the Cassini mission prevented the chugging boundary from ever being crossed. Additionally, it was possible to investigate any occurrences of actual REA-A chugging by examining chamber pressure data, albeit at an undersampled rate of only 1 Hz (vs. typical chugging frequencies for this engine around 256 Hz, near the musical note middle C). One final consumable was not tracked initially but was monitored during the final eight years of the mission—RCS thruster "starts" or deep thermal cycles. This tracking commenced with the swap to B-branch thrusters in 2009, since incurring excessive thruster starts might contribute or have contributed to thruster degradation or failure.[8] No individual Cassini RCS thruster exceeded 50% of its deep thermal cycle limit of 4000 starts, including A-branch RCS thruster operations 1997-2009 and B-branch thruster operations 2009-2017.

When Cassini entered Saturn's atmosphere, it had roughly 25% of its launch load of N_2H_4 remaining, though earlier in the mission hydrazine margins were much more precarious, particularly during periods with very frequent Reaction Wheel Assembly (RWA) biases, approaching one bias per day. This vast improvement in hydrazine margin as the mission progressed is a testament to excellent Cassini navigation, the cancellation of roughly 25% of Cassini's planned propulsive maneuvers, substantial decreases in RWA bias frequency, and exemplary performance by the spacecraft. Bipropellant margins at the moment of Saturn entry were far more perilous than monopropellant margins, 2% or less, including unusable NTO and MMH. Depending on the statistical confidence level selected, usable bipropellant was less than 1%, and in fact there was a reasonable chance of bipropellant depletion during the last few years of Cassini Solstice Mission! This will be discussed in greater detail below in the context of determining monopropellant and bipropellant usage.

Other than NTO and MMH masses and REA-A engine cycles, no other CPMS consumable exceeded 50% of its design life limits, despite two mission extensions during this already lengthy mission. One factor that helped in this regard was the swap to B-branch RCS thrusters in 2009, due to A-branch RCS thruster degradation.[8] A summary of the CPMS consumable final values at Loss of Signal (LOS) in Saturn's atmosphere on September 15, 2017 is presented below in Table 3.

Table 3. Cassini Propulsion System Consumables as of End of Mission (September 15, 2017)

CPMS Consumable	Used thru End of Mission	Lifetime Limit	% Used
NTO (Biprop Oxidizer) Mass [kg]	1833.4	1869.0	**98.1%**
MMH (Biprop Fuel) Mass [kg]	1116.1	1131.0	**98.7%**
N_2H_4 (Monoprop) Mass [kg]	98.1	132.0	74.3%
Z1A Thruster Valve Cycles [ ]	121205	273000	44.4%
Z2A Thruster Valve Cycles [ ]	106846	273000	39.1%
Z3A Thruster Valve Cycles [ ]	99874	273000	36.6%
Z4A Thruster Valve Cycles [ ]	117930	273000	43.2%
Y1A/Y3A Thruster Valve Cycles [ ]	48839	273000	17.9%
Y2A/Y4A Thruster Valve Cycles [ ]	49874	273000	18.3%
Z1A Thruster Throughput [kg]	10.43	25.0	41.7%
Z2A Thruster Throughput [kg]	8.79	25.0	35.2%
Z3A Thruster Throughput [kg]	8.94	25.0	35.8%
Z4A Thruster Throughput [kg]	10.29	25.0	41.2%
Y1A/Y3A Thruster Throughput [kg]	3.54	25.0	14.2%
Y2A/Y4A Thruster Throughput [kg]	4.49	25.0	18.0%
Z1B Thruster Valve Cycles [ ]	77728	273000	28.5%
Z2B Thruster Valve Cycles [ ]	61994	273000	22.7%
Z3B Thruster Valve Cycles [ ]	58283	273000	21.3%
Z4B Thruster Valve Cycles [ ]	76470	273000	28.0%
Y1B/Y3B Thruster Valve Cycles [ ]	52894	273000	19.4%
Y2B/Y4B Thruster Valve Cycles [ ]	58994	273000	21.6%
Z1B Thruster Throughput [kg]	8.14	25.0	32.6%
Z2B Thruster Throughput [kg]	6.39	25.0	25.6%
Z3B Thruster Throughput [kg]	5.97	25.0	23.9%
Z4B Thruster Throughput [kg]	7.73	25.0	30.9%
Y1B/Y3B Thruster Throughput [kg]	3.56	25.0	14.2%
Y1B/Y3B Thruster Throughput [kg]	3.56	25.0	14.2%
Latch Valve LV10 Cycles [ ]	10	5000	0.2%
Latch Valve LV11 Cycles [ ]	0	5000	0.0%
Latch Valve LV20 Cycles [ ]	5	5000	0.1%
Latch Valve LV21 Cycles [ ]	183	5000	3.7%
Latch Valve LV22 Cycles [ ]	2	5000	<0.1%
Latch Valve LV30 Cycles [ ]	10	5000	0.2%
Latch Valve LV31 Cycles [ ]	183	5000	3.7%
Latch Valve LV32 Cycles [ ]	2	5000	<0.1%
Latch Valve LV40 Cycles [ ]	2	5000	<0.1%
Latch Valve LV41 Cycles [ ]	1	5000	<0.1%
REA-A Cycles [ ]	183	200	**91.5%**
REA-B Cycles [ ]	0	200	0.0%

V. TCM & OTM Performance

Cassini performed 360 propulsive maneuvers over twenty years, an average of one maneuver every three weeks. This maneuver count is an order of magnitude higher than typical planetary missions, largely due to the lengthy mission duration and the relatively short orbital periods of the Cassini orbiter during its thirteen years at Saturn, often of the same order as Titan's orbital period, roughly 16 days. For many orbits of Saturn, three OTMs were planned—(1) a clean-up maneuver to correct prior satellite flyby errors, (2) an apochrone burn to optimally target the next satellite encounter, and (3) an approach maneuver to clean up apochrone burn execution errors. Roughly 25% of Cassini's planned propulsive maneuvers were able to be canceled during the mission, largely due to excellent navigation and lower-than-expected TCM and OTM execution errors. The split between main-engine and RCS maneuvers was nearly even, with 183 main-engine burns (16 TCMs and 167 OTMs) and 177 RCS maneuvers (3 TCMs and 174 OTMs) performed. Early in the mission, maneuvers with a required ΔV less than 0.5 m/s were performed on RCS thrusters, but this "crossover" point was reduced to 0.25 m/s early in the orbital tour of Saturn, given concerns about eventual hydrazine depletion at that time.

During the 6.7-year cruise period between launch and SOI, only 19 TCMs were executed, or about three per year. Therefore, OTM frequency was very high, with the remaining 341 maneuvers performed over 13.2 years, or an average of two OTMs per month. Maneuver performance for most TCMs was included in Ref. 6, but for completeness, Table 4 below finalizes this table from the reference, adding in the remaining few TCMs executed 2002-2004, including SOI. Propulsive maneuver performance throughout cruise was nominal for both main-engine and RCS TCMs. Main-engine burns used an accelerometer for burn termination (except TCM-19b, which tested burn termination on "energy" as required for SOI restarts), while RCS TCMs were truly "open loop" and thus had to be carefully predicted and executed to maintain the trajectory in Figure 1. During cruise, TCMs 3, 4, 8, 15, 16, and 22 were able to be canceled, again due to excellent navigation. RCS TCM performance was superb, within a few percent of nominal, and main-engine performance was even better, thanks to Cassini's highly accurate accelerometer.

Table 4. Cassini Trajectory Correction Maneuver (TCM) Summary Table

TCM	TCM Date [M/D/Y]	Man. Type	ΔV [m/s]	Mag. Err. [%]	Pt. Err. [°]	Primary Purpose of TCM
1	11/9/97	ME	2.7	1.70	0.63	Correct launch dispersions; validate ME TCM design
2	2/25/98	RCS	0.179	-3.51	0.51	Correct TCM-1 errors & target Venus-1; validate RCS TCM
5	12/3/98	ME	450.0	0.05	0.89	DSM: Plane-change maneuver for Venus-2; validate SOI
6	2/4/99	ME	11.552	0.19	0.07	Correct DSM (TCM-5) dispersions
7	5/18/99	RCS	0.225	-2.36	1.34	Correct TCM-6 dispersions
9	7/6/99	ME	43.485	-0.14	0.12	Correct Venus-2 flyby dispersions; Earth "walk-in" strategy
10	7/19/99	ME	5.130	-0.05	0.11	"Walk-in" maneuver for Earth flyby
11	8/2/99	ME	36.288	-0.06	0.11	"Walk-in" maneuver for Earth flyby
12	8/11/99	ME	12.246	-0.09	0.09	"Walk-in" maneuver for Earth flyby
13	8/31/99	ME	6.685	-0.16	0.31	Correct Earth flyby dispersions
14	6/14/00	ME	0.55	2.89	1.20	Flush REA-A of NTO reaction products
17	2/28/01	ME	0.51	0.13	1.20	Flush REA-A of NTO reaction products
18	4/3/02	ME	0.90	-0.44	1.0	Flush REA-A of NTO reaction products
19	5/1/03	ME	1.60	-0.06	1.0	Flush REA-A of NTO reaction products; new TCM block c/o
19a	9/10/03	RCS	0.123	2.50	0.15	New TCM block & flight S/W; 1st RCS mag boom deployed
19b	10/2/03	ME	2.02	1.10	1.01	SOI test of accelerometer burn termination on energy, not time
20	5/27/04	ME	34.71	-0.06	0.88	Biprop pressurization system c/o before SOI; Phoebe targeting
21	6/16/04	ME	3.70	-0.22	0.92	SOI arrival initial condition targeting after Phoebe flyby
--	7/1/04	ME	626.17	0.10	< 0.2	**SOI**: Orbit capture at Saturn (1st ever spacecraft to do so)

TCM-18 was the first maneuver executed after Ref. 6, and like the prior two TCMs, it was a standard flushing burn. Some CPMS components exposed to NTO unavoidably used stainless steel; therefore, a main-engine burn of at least five seconds was required at least every 400 days throughout the mission to flush out "flow decay" (iron nitrate) contaminants. TCMs 19, 19a, 19b, and 20 were executed largely to check out new functionality required for SOI and the prime mission at Saturn, including new TCM/OTM "blocks" (standard sequences of maneuver commands and timing), new flight software, and a modified burn-termination algorithm that would allow SOI burn restarts under fault conditions.[9] TCM-20 in particular was critical, the first verification of bipropellant pressurization system operation (LV10, PR1, LV20/30, CV20-23, CV30-33, etc.) in over 5.5 years. Thankfully, TCM-20 was completely nominal, setting the stage for a flawless SOI burn a mere five weeks later. LV10 was opened 108 seconds before TCM-20 start, a calculated optimal value which maximized observable regulation time for PR1 during the burn.

Following the cancellation of OTM-1, the final regulated burn of the mission took place on August 23, 2004, the Periapse Raise Maneuver (PRM), or OTM-2. Over the next thirteen years, 166 more REA-A burns were executed successfully in blowdown, with two fuel-side repressurizations provided to avoid crossing a "chugging boundary" conservatively drawn within the REA operating box (see Figure 9). Even in the presence of gross PR1 leakage, a "doubled-nested" optimization of selecting LV10 open durations for these two repressurizations was performed to maximize mission ΔV at bipropellant depletion. Extrapolating hypothetical burns to depletion starting with the final NTO and MMH mass conditions at Saturn plunge, at the moment of liquid MMH depletion, there would have been < 1 kg of remaining liquid NTO out of a launch load of 1869 kg! Bipropellant consumption uncertainties were much larger than this, of course, but this immediately demonstrates LV10 timing optimization for fuel-side repressurizations during tour was highly successful. A summary of all main-engine OTMs is provided below in Tables 5-8, including OTM number, date, actual burn time, ΔV achieved, burn ΔV magnitude error, and the principal objective of the maneuver with respect to Cassini's orbital tour at Saturn. Similar tables for RCS OTMs will be provided as well.

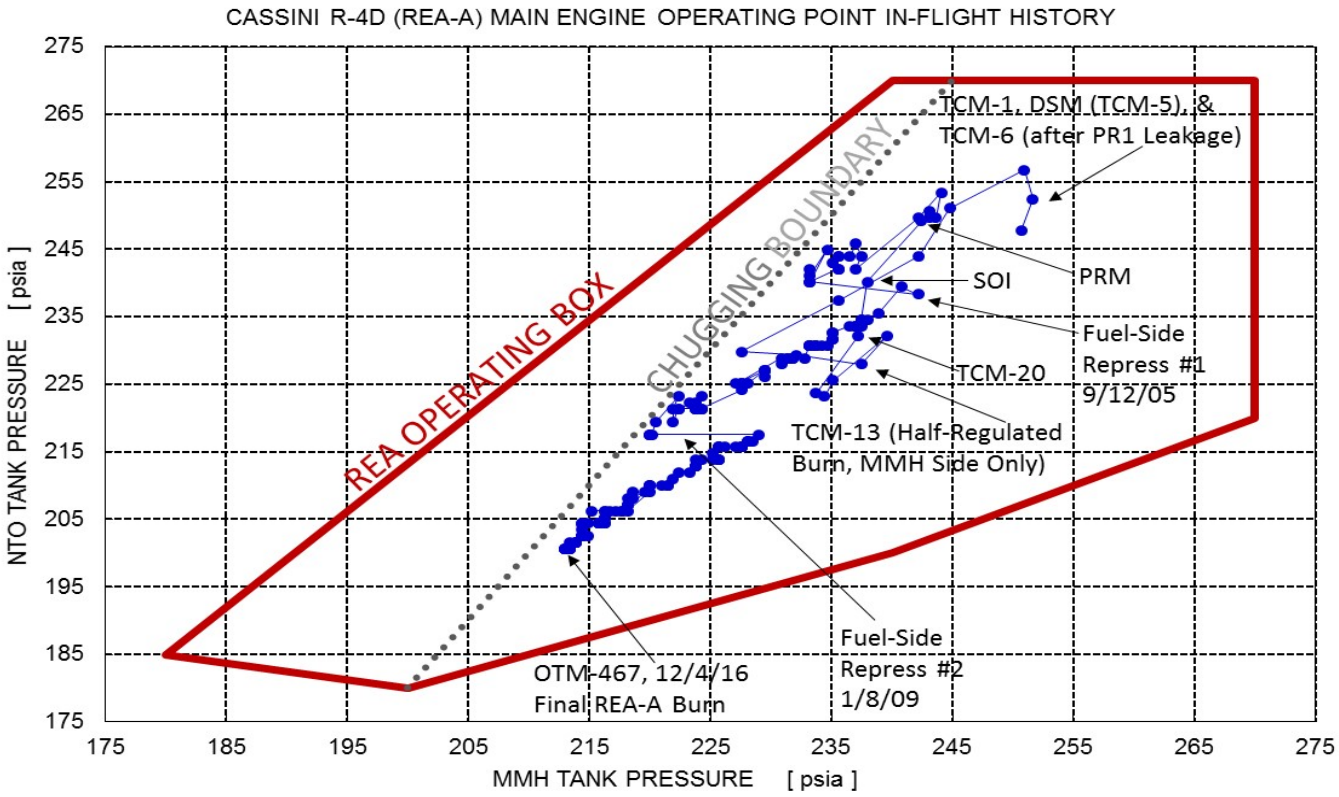

Figure 9: Cassini As-Flown Operating Point Conditions in R-4D Operating Box (in Tank Pressure Coordinates)

Table 5. Cassini ME OTM Summary Table: OTM-002 thru OTM-102 (REA-A OTMs 8/23/04 thru 4/3/07)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
2	8/23/04	3068.0	393.05	0.03	**PRM** (Periapsis Raise Maneuver); Cassini final regulated burn
3	9/7/04	3.475	0.5117	1.03	Periapsis Raise Maneuver clean-up burn; target Ta flyby
5	10/29/04	4.600	0.6456	-1.37	Clean-up burn after Ta flyby @ 1174 km altitude on 10/26/04
6	11/21/04	3.080	0.4196	0.00	Ta-to-Tb apoapsis targeting burn for 1192 km Tb flyby
8	12/17/04	84.814	11.929	-0.07	Probe Targeting Maneuver (PTM); S/C on Titan impact path
10	12/28/04	153.413	23.793	0.03	Orbiter Deflection Maneuver (ODM); OTM targeting Tc flyby
11	1/16/05	140.330	21.633	0.05	T3 targeting maneuver post Tc flyby/Huygens probe relay
12	1/28/05	119.171	18.710	0.04	Tc-to-T3 apoapsis targeting burn for T3 flyby @ 1579 km
14	2/18/05	4.458	0.7163	-0.84	T3 flyby clean-up burn targeting 3/9/05 E1 flyby @ 500 km
15	3/2/05	40.165	6.2629	0.07	T3-to-E1 apoapsis targeting burn for E1 flyby @ 500 km
17	3/12/05	2.697	0.4492	-0.48	Clean-up burn after E1 flyby @ 500 km altitude on 3/9/05
18	3/19/05	10.262	1.6200	-0.21	E1-to-T4 apoapsis targeting burn for T4 flyby @ 2404 km
20	4/4/05	5.789	0.9188	-0.89	Clean-up burn after T4 flyby @ 2404 km altitude on 3/31/05
21	4/10/05	37.350	5.8630	-0.13	T4-to-T5 apoapsis targeting burn for T5 flyby @ 1027 km
24	4/29/05	131.534	20.587	0.09	T5-to-E2 1st apoapsis targeting burn for E2 flyby on 7/14/05
25	7/8/05	2.100	0.3659	-1.75	Approach burn for E2 flyby on 7/14/05 @ 168 km altitude
26	8/3/05	16.570	2.6217	-0.24	Clean-up burn after E2 flyby @ 168 km altitude on 7/14/05
27	8/10/05	15.386	2.4164	-0.07	E2-to-T6 apoapsis targeting burn for T6 flyby @ 3660 km
29	8/25/05	9.173	1.4533	-0.41	Clean-up burn after T6 flyby @ 3660 km altitude on 8/22/05
30	8/30/05	91.350	14.356	0.04	T6-to-T7 apoapsis targeting burn for T7 flyby @ 1075 km
33	9/19/05	176.177	27.930	0.07	T7-to-Tethys1 apo targeting burn for Tethys1 flyby @ 1498km
38	10/12/05	92.725	14.832	0.02	Dione-1 clean-up burn targeting T8 flyby @ 1353 km
41	10/31/05	77.649	12.423	0.06	Clean-up burn after T8 flyby @ 1353 km altitude on 10/28/05
42	11/13/05	13.171	2.1262	-0.09	T8-to-R1 apoapsis targeting burn for Rhea-1 flyby @ 502 km
56	3/22/06	2.715	0.4706	0.76	Clean-up burn after T12 flyby @ 1949 km altitude on 3/19/06
57	4/6/06	2.206	0.3673	-0.83	T12-to-T13 apoapsis targeting burn for T13 flyby @ 1856 km
59	5/4/06	2.973	0.5101	0.98	Clean-up burn after T13 flyby @ 1856 km altitude on 4/30/06
63	6/7/06	11.975	1.9160	-0.34	T14-to-T15 apoapsis targeting burn for T15 flyby @ 1906 km
69	8/1/06	33.804	5.4918	1.25	T16-to-T17 apoapsis targeting burn for T17 flyby @ 1000 km
71	9/10/06	40.887	6.5832	0.14	Clean-up burn after T17 flyby @ 1000 km altitude on 9/7/06
72	9/14/06	50.806	8.1603	-0.10	T17-to-T18 apoapsis targeting burn for T18 flyby @ 950 km
75	10/1/06	40.25	6.4738	-0.09	T18-to-T19 apoapsis targeting burn for T19 flyby @ 980 km
78	10/17/06	5.225	0.8476	-1.40	T19-to-T20 apoapsis targeting burn for T20 flyby @ 1029 km
80	11/9/06	22.703	3.6633	-0.22	Clean-up burn after T20 flyby @ 1029 km altitude on 10/25/06
83	12/15/06	4.758	0.7914	-1.38	Clean-up burn after T21 flyby @ 1000 km altitude on 12/12/06
84	12/20/06	42.725	6.8667	-0.11	T21-to-T22 apoapsis targeting burn for T22 flyby @ 1297 km
86	12/31/06	2.823	0.4952	0.66	Clean-up burn after T22 flyby @ 1297 km altitude on 12/28/06
87	1/5/07	10.172	1.6490	-0.50	T22-to-T23 apoapsis targeting burn for T23 flyby @ 1000 km
90	1/21/07	14.694	2.3947	0.02	T23-to-T24 apoapsis targeting burn for T24 flyby @ 2631 km
93	2/7/07	1.523	0.2781	-1.46	T24-to-T25 apoapsis targeting burn for T25 flyby @ 1000 km
96	3/2/07	4.079	0.6645	-0.10	T25-to-T26 apoapsis targeting burn for T26 flyby @ 981 km
98	3/13/07	6.639	1.0747	-0.17	Clean-up burn after T26 flyby @ 981 km altitude on 3/10/07
99	3/18/07	9.955	1.6117	-0.44	T26-to-T27 apoapsis targeting burn for T27 flyby @ 1010 km
101	3/28/07	3.221	0.5402	1.83	Clean-up burn after T27 flyby @ 1010 km altitude on 3/26/07
102	4/3/07	16.712	2.6939	-0.29	T27-to-T28 apoapsis targeting burn for T28 flyby @ 991 km

Table 6. Cassini ME OTM Summary Table: OTM-105 thru OTM-195 (REA-A OTMs 4/4/07 thru 5/14/09)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
105	4/19/07	21.888	3.5325	-0.22	T28-to-T29 apoapsis targeting burn for T29 flyby @ 981 km
108	5/4/07	34.508	5.5829	-0.05	T29-to-T30 apoapsis targeting burn for T30 flyby @ 959 km
111	5/21/07	34.257	5.5339	-0.02	T30-to-T31 apoapsis targeting burn for T31 flyby @ 2299 km
113	6/1/07	4.303	0.6985	-0.20	Clean-up burn after T31 flyby @ 2299 km altitude on 5/28/07
114	6/5/07	75.567	12.237	0.00	T31-to-T32 apoapsis targeting burn for T32 flyby @ 965 km
116	6/16/07	4.585	0.7461	-1.07	Clean-up burn after T32 flyby @ 965 km altitude on 6/13/07
117	6/21/07	49.037	7.9704	0.03	T32-to-T33 apoapsis targeting burn for T33 flyby @ 1933 km
123 b/u	8/6/07	2.598	0.4268	-0.88	T33-to-T34 apoapsis burn; b/u executed due to DSN problems
125	9/2/07	2.975	0.4879	0.21	Clean-up burn after T35 flyby @ 3324 km altitude on 8/31/07
128	9/13/07	83.415	13.470	-0.09	Clean-up burn after Iapetus-1 flyby @ 1622 km on 9/10/07
131	10/5/07	8.168	1.3270	-0.30	Clean-up burn after T36 flyby @ 973 km altitude on 10/2/07
132	11/1/07	6.013	0.9770	-0.50	T36-to-T37 apoapsis targeting burn for T37 flyby @ 999 km
134	11/22/07	7.207	1.1727	0.08	Clean-up burn after T37 flyby @ 999 km altitude on 11/19/07
135	11/27/07	96.475	15.763	-0.01	T37-to-T38 apoapsis targeting burn for T38 flyby @ 1298 km
137	12/8/07	4.178	0.6809	0.15	Clean-up burn after T38 flyby @ 1298 km altitude on 12/5/07
138	12/13/07	59.073	9.6368	-0.06	T38-to-T39 apoapsis targeting burn for T39 flyby @ 970 km
141	12/29/07	12.551	2.0468	-0.27	T39-to-T40 apoapsis targeting burn for T40 flyby @ 1014 km
143	1/16/08	17.638	2.8786	-0.07	Clean-up burn after T40 flyby @ 1014 km altitude on 1/5/08
144	2/6/08	227.845	37.406	0.03	T40-to-T41 apoapsis targeting burn for T41 flyby @ 1000 km
145	2/19/08	1.725	0.2911	-2.50	Approach burn for T41 flyby on 2/22/08 @ 1000 km altitude
146	3/1/08	42.600	7.0206	-0.10	Clean-up burn after T41 flyby @ 1000 km altitude on 2/22/08
147	3/7/08	6.798	1.1240	0.27	T41-to-E3 apoapsis targeting burn for E3 flyby @ 48 km
149	3/13/08	16.673	2.7529	-0.26	Clean-up burn after E3 flyby @ 48 km altitude on 3/12/08
152	4/11/08	20.163	3.3273	-0.04	Clean-up burn after T42 flyby @ 1000 km altitude on 3/25/08
153	4/26/08	3.062	0.5153	0.60	T42-to-T43 apoapsis targeting burn for T43 flyby @ 1000 km
155	5/17/08	7.100	1.1728	-0.29	Clean-up burn after T43 flyby @ 1000 km altitude on 5/12/08
159	6/23/08	73547	12.179	-0.05	T45 (1591 km) targeting burn for **Cassini Equinox Mission**
162	8/3/08	15.401	2.5391	-0.06	Targeting maneuver for E4 flyby @ 49 km on 8/11/08
164	8/23/08	81.175	13.518	-0.07	Clean-up burn after E4 flyby on 8/11/08 @ 49 km altitude
164a	9/20/08	5.296	0.8813	-1.34	Targeting maneuver (after solar conjunction) for E5 close flyby
165	10/2/08	23.709	3.9348	-0.06	Targeting maneuver for E5 flyby @ 25 km on 10/9/08
167	10/12/08	20.039	3.3367	-0.09	Clean-up burn after E5 flyby on 10/9/08 @ 25 km altitude
168	10/17/08	41.725	6.9881	-0.07	E5-to-E6 apoapsis targeting burn for E6 flyby @ 169 km
170	11/8/08	54.600	9.0973	-0.03	Clean-up burn after T46 flyby on 11/3/08 @ 1105 km altitude
171	11/12/08	30.850	5.1495	-0.11	T46-to-T47 apoapsis targeting burn for T47 flyby @ 1023 km
173	11/23/08	4.660	0.7789	-1.03	Clean-up burn after T47 flyby on 11/19/08 @ 1023 km altitude
176	12/9/08	18.225	3.0339	-0.14	Clean-up burn after T48 flyby on 12/5/08 @ 961 km altitude
177	12/13/08	9.725	1.6197	-0.48	T48-to-T49 apoapsis targeting burn for T49 flyby @ 971 km
180	1/24/09	27.796	4.6699	-0.06	T49-to-T50 apoapsis targeting burn for T50 flyby @ 967 km
182	2/10/09	2.100	0.3637	-1.81	Clean-up burn after T50 flyby on 2/7/09 @ 967 km altitude
183	3/9/09	29.875	5.0226	-0.07	T50-to-T51 apoapsis targeting burn for T51 flyby @ 963 km
186	3/29/09	4.437	0.7482	-0.61	Clean-up burn after T51 flyby on 3/27/09 @ 963 km altitude
189	4/12/09	42.225	7.1237	-0.06	T52-to-T53 apoapsis targeting burn for T53 flyby @ 3599 km
192	4/28/09	14.817	2.4860	-0.21	T53-to-T54 apoapsis targeting burn for T54 flyby @ 3242 km
195	5/14/09	13.182	2.2224	-0.29	T54-to-T55 apoapsis targeting burn for T54 flyby @ 966 km

Table 7. Cassini ME OTM Summary Table: OTM-198 thru OTM-335 (REA-A OTMs 5/15/09 thru 11/17/12)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
198	5/30/09	8.675	1.4630	-0.40	T55-to-T56 apoapsis targeting burn for T56 flyby @ 968 km
200	6/10/09	12.674	2.1396	-0.28	Clean-up burn after T56 flyby on 6/6/09 @ 968 km altitude
203	6/26/09	14.335	2.4245	-0.01	Clean-up burn after T57 flyby on 6/22/09 @ 955 km altitude
206	7/12/09	20.750	3.5129	-0.14	Clean-up burn after T58 flyby on 7/8/09 @ 966 km altitude
209	7/28/09	37.080	6.2945	-0.02	Clean-up burn after T59 flyby on 7/24/09 @ 956 km altitude
213	8/16/09	76.475	13.002	-0.02	T60-to-T61 apoapsis targeting burn for T61 flyby @ 961 km
215	8/29/09	2.973	0.5149	0.25	Clean-up burn after T61 flyby on 8/25/09 @ 961 km altitude
216	9/5/09	26.291	4.4757	-0.10	T61-to-T62 apoapsis targeting burn for T62 flyby @ 1300 km
218	10/16/09	4.969	0.8441	-0.87	Clean-up burn after T62 flyby on 10/12/09 @ 1300 km altitude
219	10/21/09	24.416	4.1621	-0.16	T62-to-E7 apoapsis targeting burn for E7 flyby @ 99 km
221	11/5/09	1.725	0.3026	-3.01	Clean-up burn after E7 flyby on 11/2/09 @ 99 km altitude
224	11/22/09	14.937	2.5532	-0.13	Clean-up burn after E8 flyby on 11/21/09 @ 1597 km altitude
227	12/15/09	4.184	0.7161	-0.36	Clean-up burn after T63 flyby on 12/12/09 @ 4848 km altitude
228	12/20/09	13.065	2.2249	-0.28	T63-to-T64 apoapsis targeting burn for T64 flyby @ 951 km
231	1/5/10	47.082	8.0536	0.05	T64-to-T65 apoapsis targeting burn for T65 flyby @ 1074 km
233	1/16/10	13.327	2.2660	-0.17	Clean-up burn after T65 flyby on 1/12/10 @ 1074 km altitude
234	1/21/10	35.530	6.0749	0.05	T65-to-T66 apoapsis targeting burn for T66 flyby @ 7486 km
236	2/1/10	36.225	6.1970	-0.10	Clean-up burn after T66 flyby on 1/28/10 @ 7486 km altitude
240	3/26/10	17.527	3.0040	0.08	R2-to-T67 apoapsis targeting burn for T67 flyby @ 7438 km
242	4/10/10	52.687	9.0417	0.00	Clean-up burn after D2 flyby on 4/7/10 @ 506 km altitude
245	4/29/10	33.250	5.7160	0.03	Clean-up burn after E9 flyby on 4/28/10 @ 100 km altitude
246	5/11/10	51.545	8.8840	0.00	E9-to-E10 apoapsis targeting burn for E10 flyby @ 437 km
248	5/23/10	4.944	0.8458	-0.76	Clean-up burn after T68 flyby on 5/20/10 @ 1398 km altitude
249	5/28/10	62.428	10.768	0.01	T68-to-T69 apoapsis targeting burn for T69 flyby @ 2042 km
252	6/13/10	7.209	1.2380	-0.08	T69-to-T70 apoapsis targeting burn for T70 flyby @ **878 km**
254	6/24/10	5.026	0.8707	-0.42	Clean-up burn after T70 flyby on 6/21/10 @ 878 km altitude
255	6/30/10	36.208	6.2570	-0.01	T70-to-T71 apoapsis targeting burn for T71 flyby @ 1004 km
257	7/10/10	4.829	0.8322	-0.70	Clean-up burn after T71 flyby on 7/7/10 @ 1004 km altitude
258	7/18/10	39.144	6.7640	-0.09	T71-to-E11 apoapsis targeting burn for E11 flyby @ 2555 km
261	9/3/10	14.145	2.4380	-0.06	E11-to-T72 periapsis targeting burn for T72 flyby @ 8178 km
267	11/21/10	13.100	2.2470	-0.13	T73-to-E12 apo burn; 1st ME OTM **Cassini Solstice Mission**
275	1/14/11	15.969	2.7630	-0.11	Clean-up burn after R3 flyby on 1/11/11 @ 70 km altitude
291	9/20/11	29.419	5.0560	0.04	T78-to-E14 apoapsis targeting burn for E14 flyby @ 99 km
299	11/9/11	12.029	2.0860	-0.10	Clean-up burn after E16 flyby on 11/6/11 @ 496 km altitude
300	11/24/11	17.175	2.9750	0.00	E16-to-Dione3 periapsis targeting burn for D3 flyby @ 99 km
303	12/17/11	2.923	0.5073	0.83	Clean-up burn after T79 flyby on 12/13/11 @ 3586 km altitude
312	3/10/12	20.600	3.5645	-0.29	T82-to-E17 periapsis targeting burn for E17 flyby @ 74 km
318	4/24/12	1.319	0.2458	2.30	E18-to-E19 apoapsis targeting burn for E19 flyby @ 74 km
321	5/14/12	47.600	8.2696	-0.03	E19-to-T83 apoapsis targeting burn for T83 flyby @ 955 km
324	5/30/12	21.395	3.7149	0.02	T83-to-T84 apoapsis targeting burn for T84 flyby @ 959 km
326	6/10/12	2.350	0.4216	-0.01	Clean-up burn after T84 flyby on 6/7/12 @ 959 km altitude
327	6/21/12	58.027	10.118	-0.01	T84-to-T85 apoapsis targeting burn for T85 flyby @ 1012 km
330	8/7/12	24.973	4.3551	0.06	T85-to-T86 apoapsis targeting burn for T86 flyby @ 956 km
333	10/9/12	4.332	0.7638	0.40	T86-to-T87 apoapsis targeting burn for T86 flyby @ 973 km
335	11/17/12	1.418	0.2516	-1.24	Clean-up burn after T87 flyby on 11/13/12 @ 973 km altitude

Table 8. Cassini ME OTM Summary Table: OTM-336 thru OTM-467 (REA-A OTMs 11/18/12 thru 9/15/17)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
336	11/22/12	28.401	4.9597	0.01	T87-to-T88 apoapsis targeting burn for T88 flyby @ 1014 km
339	1/30/13	9.528	1.6538	-0.22	T88-to-T89 periapsis targeting burn for T89 flyby @ 1978 km
341	2/24/13	8.310	1.4470	-0.19	T89-to-R4 periapsis targeting burn for R4 flyby @ 997 km
342	3/2/13	1.451	0.2610	-1.26	T89-to-R4 apoapsis targeting burn for R4 flyby @ 997 km
348	4/30/13	2.817	0.4954	0.19	T90-to-T91 periapsis targeting burn for T91 flyby @ 970 km
351	6/11/13	4.675	0.8187	-0.17	T91-to-T92 periapsis targeting burn for T92 flyby @ 964 km
353 b/u	7/15/13	1.350	0.2491	-1.56	T92 clean-up burn; b/u to save N_2H_4 (large RCS → small ME)
354	7/19/13	12.942	2.2689	0.15	T92-to-T93 apoapsis targeting burn for T93 flyby @ 1400 km
357	8/7/13	20.625	3.6194	0.14	T93-to-T94 apoapsis targeting burn for T94 flyby @ 1400 km
363	11/2/13	2.000	0.3654	0.35	T95-to-T96 apoapsis targeting burn for T96 flyby @ 1400 km
366	12/17/13	2.142	0.3793	-1.71	T96-to-T97 apoapsis targeting burn for T97 flyby @ 1400 km
372	2/17/14	9.591	1.6809	-0.12	T98-to-T99 apoapsis targeting burn for T99 flyby @ 1500 km
375	3/20/14	3.046	0.5448	0.52	T99-to-T100 apoapsis targeting burn for T100 flyby @ 963 km
387	8/9/14	71.019	12.466	0.05	T103-to-T104 apoapsis target burn for T104 flyby @ 964 km
390	9/7/14	7.225	1.2616	-0.34	T104-to-T105 apoapsis target burn for T105 flyby @ 1400 km
393	10/9/14	6.034	1.0613	-0.32	T105-to-T106 apoapsis target burn for T106 flyby @ 1013 km
399	12/29/14	5.473	0.9645	-0.40	T107-to-T108 apoapsis target burn for T108 flyby @ 970 km
402	1/31/15	7.214	1.2724	0.37	T108-to-T109 apoapsis target burn for T109 flyby @ 1200 km
404	2/15/15	2.812	0.4950	-0.43	Clean-up burn after T109 flyby on 2/12/15 @ 1200 km altitude
422	10/2/15	1.375	0.2464	-2.26	Clean-up burn after T113 flyby on 9/28/15 @ 1036 km altitude
423	10/6/15	14.949	2.6204	-0.20	T113-to-E20 apoapsis target burn for E20 flyby @ 1839 km
435	12/30/15	16.974	2.9843	-0.07	E22-to-T115 apoapsis target burn for T115 flyby @ 3817 km
438	1/23/16	38.725	6.8442	-0.04	T115-to-T116 apoapsis target burn for T116 flyby @ 1400 km
440	2/3/16	3.270	0.5785	-0.72	Clean-up burn after T116 flyby on 2/1/16 @ 1400 km altitude
441	2/8/16	4.202	0.7448	-0.28	T116-to-T117 apoapsis target burn for T117 flyby @ 1018 km
444	3/25/16	44.952	7.9562	0.06	T117-to-T118 apoapsis target burn for T118 flyby @ 990 km
447	4/22/16	9.975	1.7641	-0.19	T118-to-T119 apoapsis target burn for T119 flyby @ 971 km
449	5/9/16	3.079	0.5560	0.89	Clean-up burn after T119 flyby on 5/6/16 @ 971 km altitude
452	6/11/16	1.375	0.2465	-3.05	Clean-up burn after T120 flyby on 6/7/16 @ 975 km altitude
453	7/17/16	11.466	2.0205	-0.22	T120-to-T121 apoapsis target burn for T121 flyby @ 976 km
456	8/2/16	4.461	0.7934	-0.08	T121-to-T122 apoapsis target burn for T122 flyby @ 1599 km
467	12/4/16	5.600	0.9898	-0.37	Clean-up burn after T125 flyby on 11/29/16 @ 3223 km alt.

Tables 4-8 likely have historical and archival interest, but it is difficult to assess maneuver performance trends in such a tabular format. One key performance metric is ΔV magnitude error as a function of executed ΔV; this comparison is presented below in Figure 10, on a semi-log scale. Since all but one of Cassini's 183 REA-A burns terminated on accelerometer control, any consistent bias in ΔV magnitude error should have been due to accelerometer bias, rather than propulsion errors. In fact, accelerometer scale factor was (minorly) updated a few times during the mission based on small but consistent TCM and OTM execution errors, particularly for ΔV's in the 1-7 m/s range. For very small main-engine burns, ΔV magnitude error percentage was high, as expected, due to uncertainties in burn timing and termination accounting, tail-off impulse, and minor RCS Z-thruster deadbanding ΔV around the time of the burn. This trend is evident in Figure 10, but overall Cassini REA-A performance was exemplary over two decades of operation in space. If the R-4D was good enough for Apollo astronauts to rely on this engine for attitude control on the lunar module, it was likely an excellent choice for Cassini. Perhaps the only regret (from propulsion engineers on Cassini) was that REA-B was never operated in flight, given mission safety concerns for actuating the 20-year-old pyrotechnic devices that would have been required to bring it on line near the end of mission.

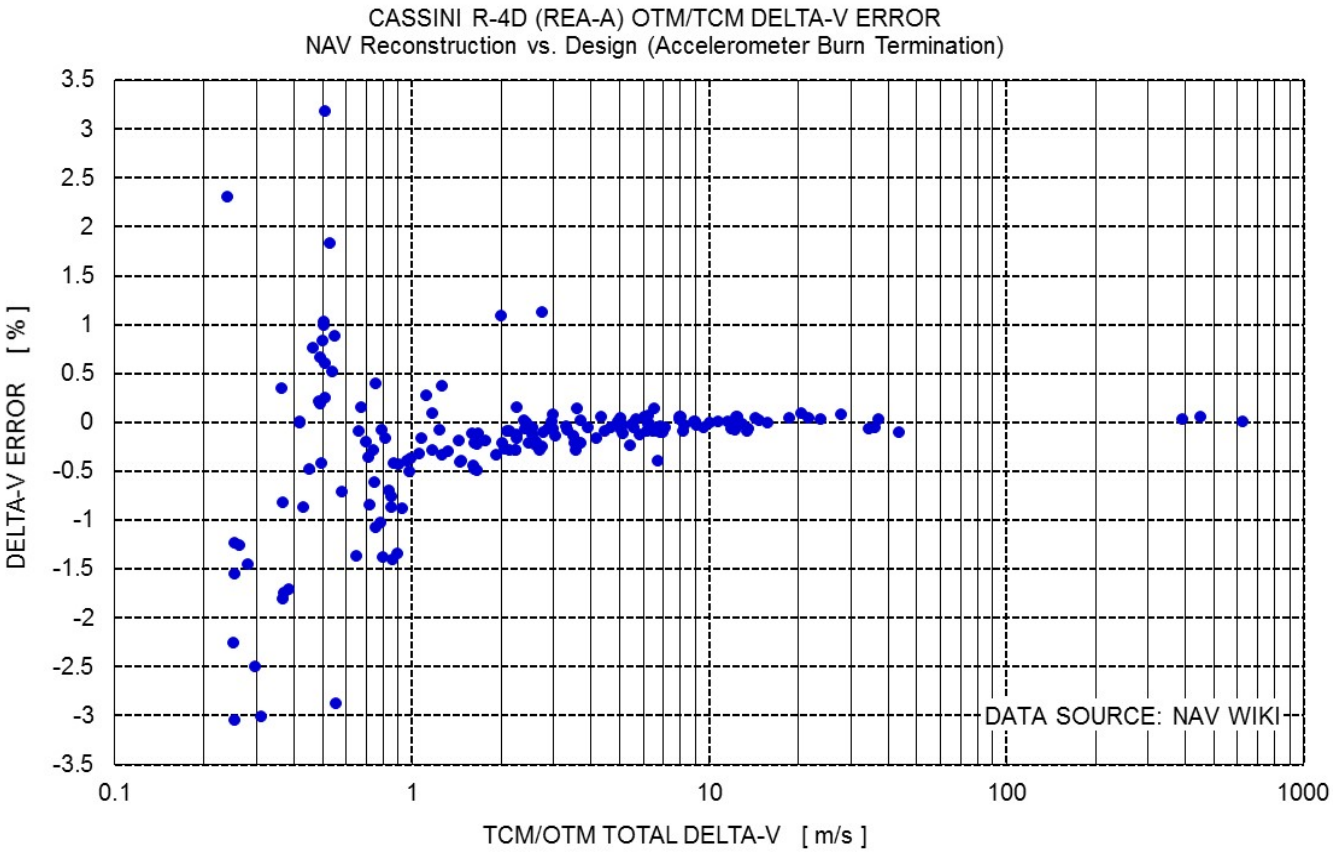

Figure 10: Cassini R-4D (REA-A) OTM/TCM ΔV Magnitude Error vs. Total ΔV Executed

Nearly all points in Figure 10 are within quite narrow error ranges, typically ± a few percent for ΔV < 1 m/s, ± 0.5% for ΔV = 1-7 m/s, and basically ± 0.1% for ΔV > 7 m/s. From left to right, the two outliers in Figure 10 with greater than 1% errors for ΔV = 2-3 m/s were TCM-19b and TCM-1, respectively. The former was the SOI energy algorithm test TCM, essentially a timed burn. As such, ΔV errors above 1% are not surprising. TCM-1 was by definition an uncalibrated burn for both the CPMS and accelerometer, so larger ΔV errors are to be expected as well.

Given the ability of the accelerometer to cover a "multitude of sins" in CPMS performance, within reasonable minimum and maximum burn time ranges, it may be more telling to track the error in predicted vs. actual burn time as a function of executed ΔV. These data are presented below in Figure 11, and indeed a consistent bias is evident in the data. Minor uncertainties in burn time for the largest three maneuvers (PRM, DSM, and SOI) were related more to pressurization system performance, but there is otherwise a clear trend of increased burn time prediction error for smaller and smaller maneuvers, especially for maneuvers with ΔV < 3 m/s. These burns were all < 20 seconds in duration, so inaccurate transient modeling in Cassini's bipropellant system prediction tool (bspt) very likely explains the trend seen in Figure 11. For nearly all Cassini REA-A TCMs and OTMs, minimum burn time was set to one second (the REA minimum burn time allowed by AACS), while the maximum burn time was specified to be the mean value +5% or +1 second, whichever was larger. These settings comfortably enveloped the trend in Figure 11, assuring burn termination by the accelerometer for all Cassini main-engine burns (again, except for TCM-19b, by design). This is particularly true since bspt nearly always overpredicted the actual burn time for main-engine maneuvers, especially for short burns. Despite the clear trend in Figure 11, there was never any need to poke "under the hood" in bspt to improve short-burn transient modeling, particularly given twenty years of highly successful operation by the accelerometer. In contrast, an accelerometer failure would have required subsequent timed main-engine burns during the mission, and great effort would have been expended to more accurately predict small maneuvers in that case.

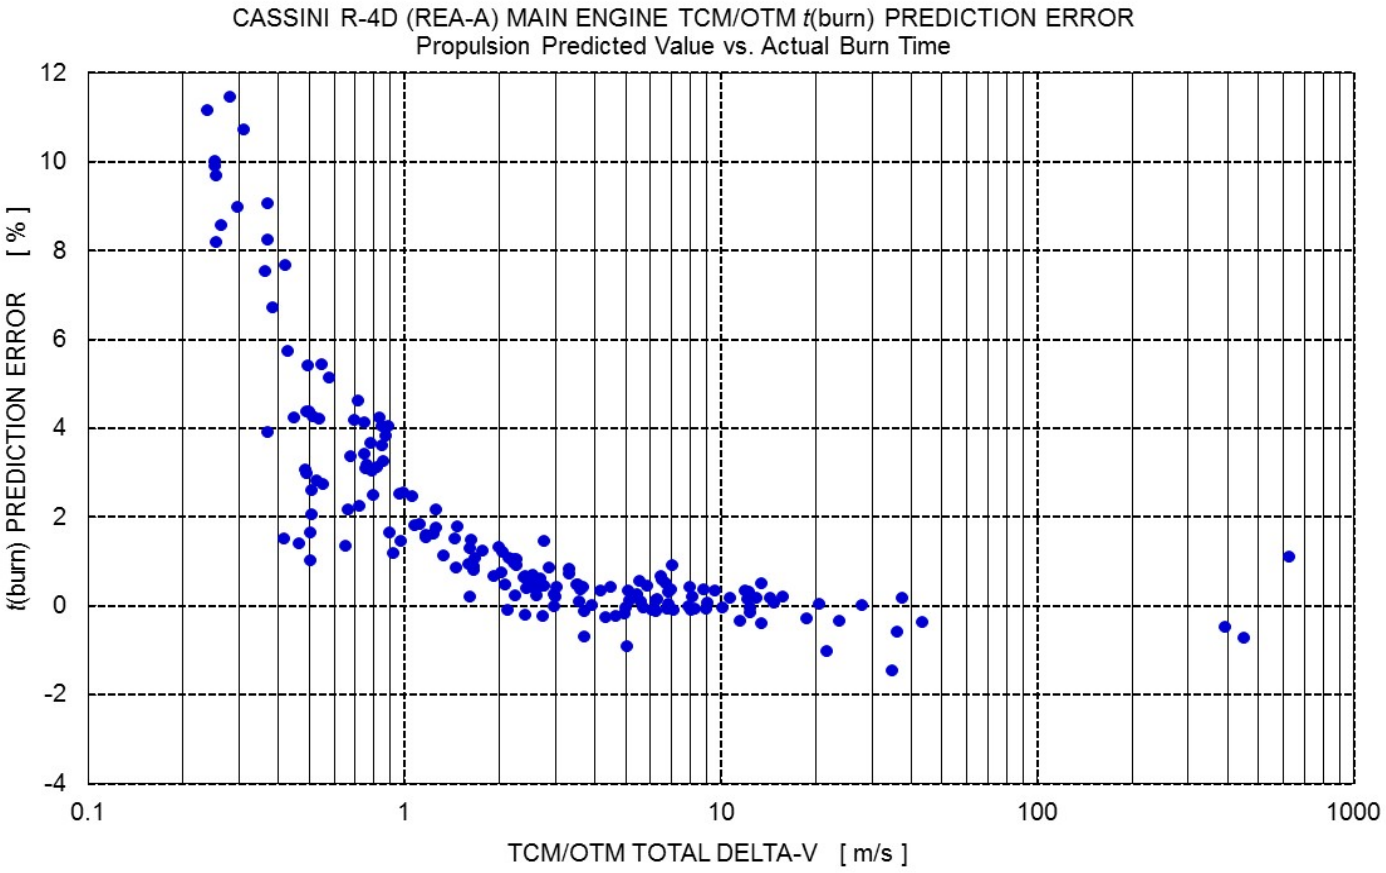

Figure 11: Cassini R-4D (REA-A) OTM/TCM Burn Time Prediction Error vs. Total ΔV Executed

One final summary plot covering most Cassini main-engine maneuvers is presented below in Figure 12, an attempt to identify and characterize any potential chamber pressure (Pc) roughness over the mission. As was mentioned previously, Pc telemetry frequency was no better than 1 Hz, quite undersampled vs. the expected R-4D chugging frequency around 256 Hz. Despite this limitation, a reasonable "proxy" for in-flight roughness was calculated by dividing the standard deviation of all in-flight 1-Hz Pc measurements (for a given burn) by the average Pc during that burn. Peak-to-peak roughness as measured during ground test is by definition higher than this in-flight derived quantity, but actual Pc roughness should have been visible in flight should it have occurred, even when undersampled at 1 Hz. Of course, when massively undersampled, true peak roughness values could easily be missed in flight.

Figure 12 includes REA-A OTMs after PRM (OTM-2) only, since this Pc roughness proxy technique was developed during tour rather than cruise. However, plots of Pc vs. burn time for all 16 main-engine TCMs (including DSM and SOI) and PRM showed miniscule Pc roughness, so including these data in Figure 12 would not change any conclusions. For burn times less than 2 seconds, only one 1-Hz sample was obtained, so the standard deviation was by definition zero. Even for much longer main-engine OTMs, some hundreds of seconds in duration, there were never any hints of chamber pressure roughness over the twenty-year Cassini mission. The data of Figure 12 would have to have been at least an order of magnitude higher before Pc roughness was suspected. This is particularly true because Pc transducers were 8-bit sensors (0-255 DN, or Data Number), so the "roughness" level observed in Figure 12 is actually equivalent to the DN uncertainty of the Pc measurements themselves (i.e., it is within noise levels).

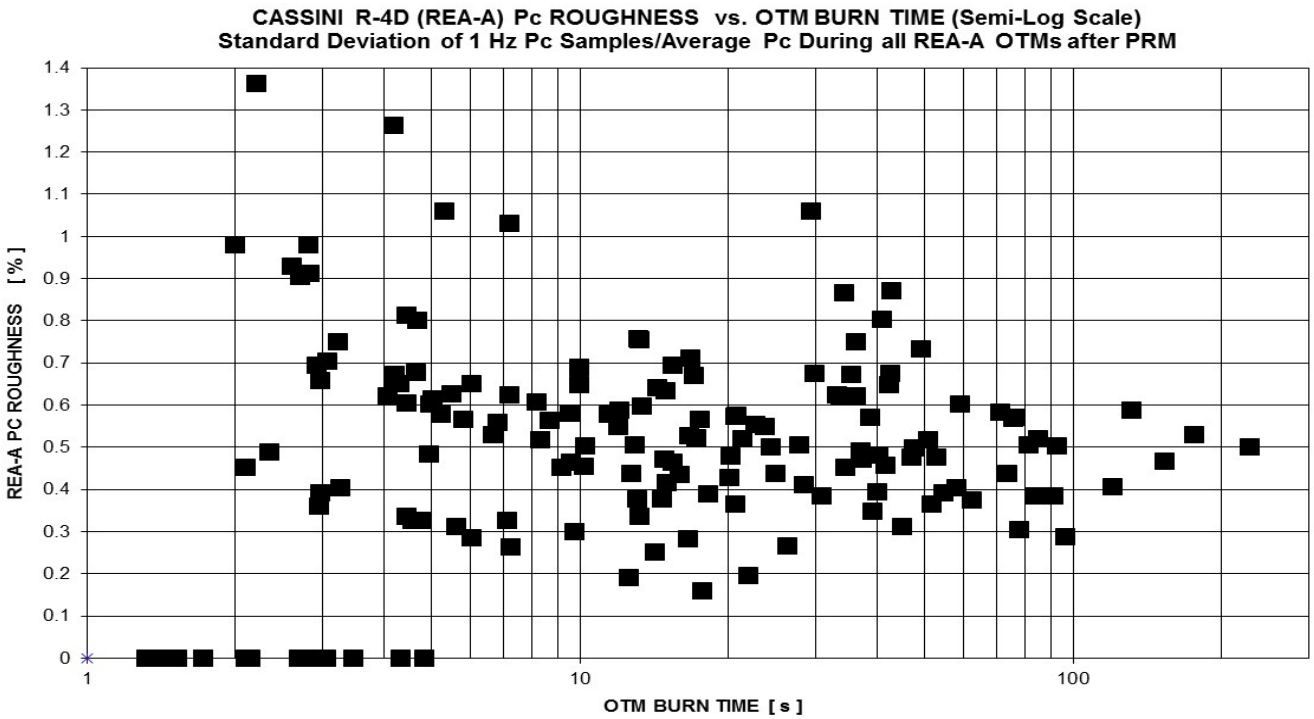

Figure 12: Cassini R-4D (REA-A) Pc Roughness Proxy (1σ Deviation of 1-Hz Samples/Average) vs. Burn Time

Before presenting summary plots and tables for RCS TCMs, for future bipropellant missions it may be useful to more completely archive maneuver performance during Cassini's five regulated maneuvers. TCM-1 was anticipated to be a regulated burn, but PR1 leakage during initial pressurization allowed this TCM to be executed in blowdown mode.[6] In fact, PR1 leakage at initial pressurization, before DSM (TCM-5), and after DSM help explain the rather convoluted path taken in Cassini's main-engine operating box (refer back to Figure 9) , especially early in the mission. Otherwise, maneuver performance during four regulated TCMs and PRM (OTM-2) was flawless. DSM was quite literally the perfect practice burn for SOI, since they were nearly the same duration. During TCM-13, a clever "half-regulated" burn (MMH-side only pressurization during the maneuver) not only prevented chugging boundary crossing, it also set up acceptable blowdown burn conditions until TCM-20 and minimized LV10 and PR1 cycles to mitigate potential leakage. As was mentioned before, TCM-20 was essentially the pressurization system check-out burn a mere five weeks before SOI, while a regulated PRM allowed just enough NTO tank repressurization to pyro-isolate the NTO side for the remainder of the mission. Following PRM, Cassini executed only blowdown main-engine maneuvers, with two fuel-side repressurizations in between OTMs to avoid breaching the chugging boundary. Pressurization system performance during all helium transfer events through LV10 and PR1 will be assessed below.

Table 9 summarizes DSM, TCM-13, TCM-20, SOI, and PRM performance, based on the example of Table 6 from Ref. 3. All values in Table 9 were as expected, even in the face of gross PR1 leakage likely due to particulate contamination. This is because PR1 helium flow rates during Cassini's five regulated burns were far higher than worst-case leakage rates through the prime regulator, and other than its sealing performance, PR1 operation in flight was flawless and identical to ground test. Table 9 parameters not previously covered in Tables 4-5 include LV10 opening/closing times, average tank pressures, PR1 outlet pressure, Pc, thrust, Isp, mixture ratio, propellant flow rates and mass quantities, and total burn impulse. All values in Table 9 were well within expected ranges based on ground test data. LV10 standard timing was opening 78 s before the burn and closing "asap" after the 7STOP command (latest possible burn termination), but TCM-13 and TCM-20 LV10 opening times were optimized to maximize PR1 regulation time, DSM LV10 closing timing enabled future blowdown TCMs, and PRM LV10 closing timing was selected to maximize tank pressurization while still accommodating a worst-case (filter-hole sized) particle at PR1.

Table 9: Summary of Cassini's Regulated Main-Engine TCMs/OTMs (DSM, TCM-13, TCM-20, SOI, & PRM)

Parameter	DSM (TCM-5)	TCM-13	TCM-20	SOI	PRM (OTM-2)
Date [D/M/Y]	12/3/98	8/31/99	5/27/04	7/1/04	8/23/04
ΔV [m/s]	450.0	6.685	34.71	626.17	393.05
t(burn) [s]	5254.884	72.100	361.896	5780.250	3067.840
LV10 Opening Time	T – 78 s	T – 168 s	T – 108 s	T – 78 s	T – 78 s
LV10 Closing Time	7STOP + 2400 s	7STOP + 14 s	7STOP + 29 s	7STOP + 31 s	7STOP + 1897 s
P(NTO, MMH) [psia]	239, 237	229, 238	238, 238	240, 238	240, 238
PREG [psia]	255	255	254	256	255
Pc(av) [psia]	98	96	98	98	98
Thrust(av) [N]	439.9	427.3	437.2	442.9	441.8
Isp(av) [s]	304.4	296.3	305.6	308.0	306.6
Mixture Ratio [ ]	1.654	1.585	1.640	1.658	1.664
∂M(NTO)/∂t [g/s]	91.82	90.15	90.63	91.45	91.76
∂M(MMH)/∂t [g/s]	55.53	56.87	55.26	55.17	55.15
∂M(prop)/∂t [g/s]	147.35	147.02	145.89	146.62	146.91
ΔM(NTO) [kg]	482.5	6.5	32.8	528.6	281.5
ΔM(MMH) [kg]	291.8	4.1	20.0	318.9	169.2
ΔM(Prop) [kg]	774.3	10.6	52.8	847.5	450.7
Impulse(tot) [Ns]	2.311e7	3.081e4	1.582e5	2.560e7	1.355e7

Turning now to RCS OTMs, Table 10 below summarizes all 45 A-branch RCS OTMs executed during the four-year prime mission and the first few months of the equinox mission. Maneuver performance was quite good for all A-branch RCS OTMs, though ΔV error percentages were often higher for small RCS burns. This is as expected, however, since the vagaries of post-burn deadband tightening and deadbanding under RCS control make it difficult to deliver ΔV accuracies better than a few mm/s. As the mission progressed, error modeling for post-burn deadband tightening ΔV and deadbanding ΔV improved, so open-loop RCS OTMs became more accurate. Ironically, OTM cancellation frequency did not increase, however, since the project became less willing to cancel maneuvers as both hydrazine and especially bipropellant margins became tight. Figure 13 is the RCS analogue to Figure 10, displaying A-branch RCS OTM ΔV magnitude error as a function of delivered ΔV. Errors were nearly always within ±3%, an excellent result for open-loop maneuvers. Other than OTM-169, outliers had very small ΔV's and/or occurred early during tour (e.g., OTM-9, with big errors in deadband modeling, off-scale high in Figure 13).

It is difficult to discern from Table 10 or even Figure 13, but OTM-169 on October 29, 2008 raised more than a few eyebrows on the Cassini project. This approach burn for the Enceladus-6 flyby had a multi-sigma underburn, one that ended up incurring a rather hefty 5-m/s penalty to correct the trajectory before Cassini's T48 flyby five weeks later. Upon further examination, it was determined the Z4A thruster and especially the Z3A thruster had pronounced Pc roughness during OTM-169, enough to affect delivered burn ΔV. Minor increases in Pc roughness were noted earlier in the mission, but they were (1) temporary, (2) not as severe, (3) concurrent on all four A-branch Z-thrusters simultaneously, and (4) perhaps related to "sining" instead of catalyst bed degradation.[8] The Pc roughness observed in OTM-169 (and, in retrospect, also in OTM-160 and OTM-166) was something else entirely. This may be seen graphically in Figure 14, which shows a proxy for Pc roughness for all A-branch RCS maneuvers by Z-thruster. Similar to the main-engine Pc roughness proxy presented in Figure 12, this quantity was also determined by dividing the standard deviation of 1-Hz samples by the average Pc value, but this time for only the first 20 seconds of the RCS burn. This was necessary because Cassini RCS TCMs/OTMs used off-pulsing for attitude control, and Pc measurements were affected by thruster valve cycling for engines connected via the same manifold. Before Huygens probe release on Christmas Eve, 2004, all RCS TCMs and OTMs had 100% duty cycle for Z3A. After probe release, spacecraft mass and inertial property changes caused the Z4A thruster to fire continuously during all subsequent A-branch RCS OTMs. It is quite interesting the two most severely degraded Cassini thrusters were the ones which were used in steady-state conditions during the mission, but this may only be coincidence. Interested readers may explore Ref. 8 in detail to learn more about Cassini A-branch MR-103H thruster degradation.

Table 10. Cassini RCS OTM Summary Table: OTM-004 thru OTM-178 (All A-Branch RCS OTMs)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
4	10/23/04	464.375	0.3831	2.90	Approach maneuver for Ta flyby on 10/26/04 @ 1174 km
9	12/23/04	18.875	0.0206	16.89	Probe Targeting Maneuver (PTM) clean-up OTM for probe rls.
10a	1/3/05	147.625	0.1388	3.04	Probe relay final target burn for 60000 km Tc flyby on 1/14/15
13	2/12/05	220.500	0.2076	0.88	Approach burn for T3 flyby on 2/15/05 @ 1579 km altitude
22	4/14/05	67.625	0.0648	1.09	Approach burn for T5 flyby on 4/16/05 @ 1027 km altitude
31	9/3/05	66.875	0.0646	2.38	Approach burn for T7 flyby on 9/7/05 @ 1075 km altitude
35	9/28/05	321.250	0.2963	0.51	Clean-up burn after Tethys-1/Hyperion-1 flybys targeting D1
39	10/21/05	96.625	0.0914	0.99	Dione1-to-T8 apoapsis targeting burn for T8 flyby @ 1353 km
43	11/23/05	63.125	0.0604	0.17	Approach burn for Rhea-1 flyby on 11/26/05 @502 km altitude
44	11/28/05	262.625	0.2409	1.43	Clean-up burn after Rhea-1 flyby targeting T9 flyby
47	12/30/05	198.750	0.1816	-0.66	Clean-up burn after T9 flyby on 12/26/05 @ 10411 km altitude
51	2/2/06	203.125	0.1852	-0.54	T10-to-T11 apoapsis targeting burn for T11 flyby @ 1812 km
53	3/2/06	291.500	0.2634	-0.57	Clean-up burn after T11 flyby targeting T12 flyby @ 1949 km
58	4/27/06	53.125	0.0786	4.11	Approach burn for T13 flyby on 4/30/06 @ 1856 km altitude
61	5/18/06	85.125	0.1221	3.12	Approach burn for T14 flyby on 5/20/06 @ 1879 km altitude
64	6/28/06	47.750	0.0694	2.21	Approach burn for T15 flyby on 7/2/06 @ 1906 km altitude
65	7/5/06	96.875	0.1378	0.29	Clean-up burn after T15 flyby targeting T16 flyby @ 950 km
70	9/4/06	164.125	0.2275	-0.09	Approach burn for T17 flyby on 9/7/06 @ 1000 km altitude
76	10/6/06	26.500	0.0410	2.50	Approach burn for T19 flyby on 10/9/06 @ 980 km altitude
79	10/22/06	44.250	0.0637	2.74	Approach burn for T20 flyby on 10/25/06 @ 1029 km altitude
81 b/u	11/27/06	160.125	0.2203	0.23	T20-to-T21 apoapsis burn; b/u executed to save 0.6 m/s ΔV
88	1/10/07	27.625	0.0419	0.48	Approach burn for T23 flyby on 1/13/07 @ 1000 km altitude
89	1/16/07	158.500	0.2131	-0.19	Clean-up burn after T23 flyby targeting T24 flyby @ 2631 km
91	1/26/07	7.625	0.0150	-1.96	Approach burn for T24 flyby on 1/29/07 @ 2631 km altitude
94	2/19/07	28.000	0.0414	-0.24	Approach burn for T25 flyby on 2/22/07 @ 1000 km altitude
100	3/22/07	49.750	0.0695	1.16	Approach burn for T27 flyby on 3/26/07 @ 1010 km altitude
103	4/7/07	25.875	0.0378	0.80	Approach burn for T28 flyby on 4/10/07 @ 991 km altitude
106	4/24/07	8.875	0.0171	8.92	Approach burn for T29 flyby on 4/26/07 @ 981 km altitude
109	5/9/07	14.500	0.0253	3.69	Approach burn for T30 flyby on 5/12/07 @ 959 km altitude
115	6/11/07	25.500	0.0364	-2.15	Approach burn for T32 flyby on 6/13/07 @ 965 km altitude
118	6/26/07	6.875	0.0132	-1.49	Approach burn for T33 flyby on 6/29/07 @ 1933 km altitude
119	7/3/07	13.875	0.0243	6.11	Clean-up burn after T33 flyby targeting T34 flyby @ 1332 km
121	7/15/07	7.000	0.0137	1.48	Approach burn for T34 flyby on 7/19/07 @ 1332 km altitude
129	9/17/07	79.000	0.1029	0.10	Clean-up burn after Iapetus-1 flyby targeting T36 flyby 973 km
130	9/28/07	14.375	0.0242	2.98	Approach burn for T36 flyby on 10/2/07 @ 973 km altitude
133	11/15/07	50.750	0.0675	1.05	Approach burn for T37 flyby on 11/19/07 @ 999 km altitude
136	12/2/07	10.625	0.0195	5.98	Approach burn for T38 flyby on 12/5/07 @ 1298 km altitude
139	12/18/07	7.125	0.0137	0.74	Approach burn for T39 flyby on 12/20/07 @ 970 km altitude
150	3/18/08	40.500	0.0552	2.22	E3-to-T42 apoapsis targeting burn for T42 flyby @ 1000 km
156	5/22/08	155.125	0.1958	-0.26	T43-to-T44 apoapsis targeting burn for T44 flyby @ 1400 km
160	7/27/08	139.000	0.1695	-2.02	Approach burn for T45; 1st OTM **Cassini Equinox Mission**
166	10/6/08	8.125	0.0150	2.86	Approach burn for E5 flyby on 10/9/08 @ 25 km
169	10/29/08	191.000	0.2248	-3.13	Approach burn for E6 flyby on 10/31/08 @ 169 km; Pc rough
175	12/1/08	57.500	0.0683	-0.10	Approach burn for T48 flyby on 12/5/08 @ 961 km; Pc rough
178	12/17/08	19.625	0.0270	2.39	Approach burn for T49 flyby on 12/21/08 @ 971 km; Pc rough

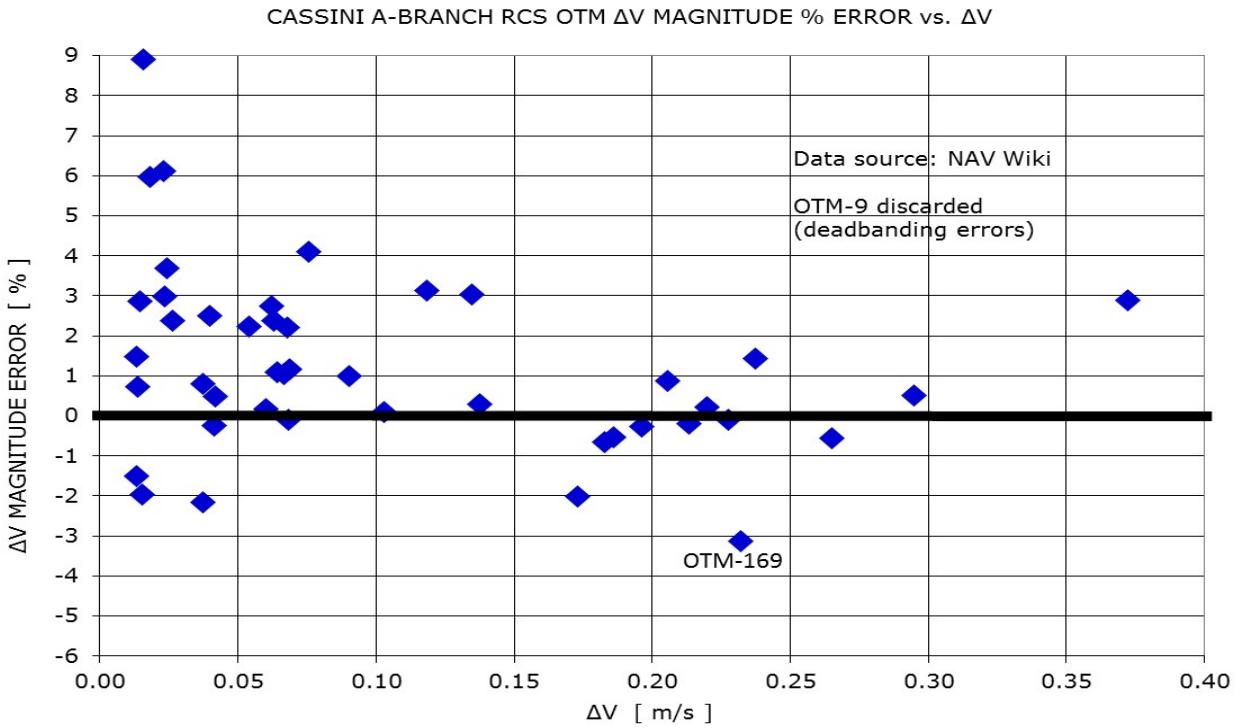

Figure 13: Cassini A-Branch RCS OTM ΔV Magnitude % Error vs. Executed ΔV

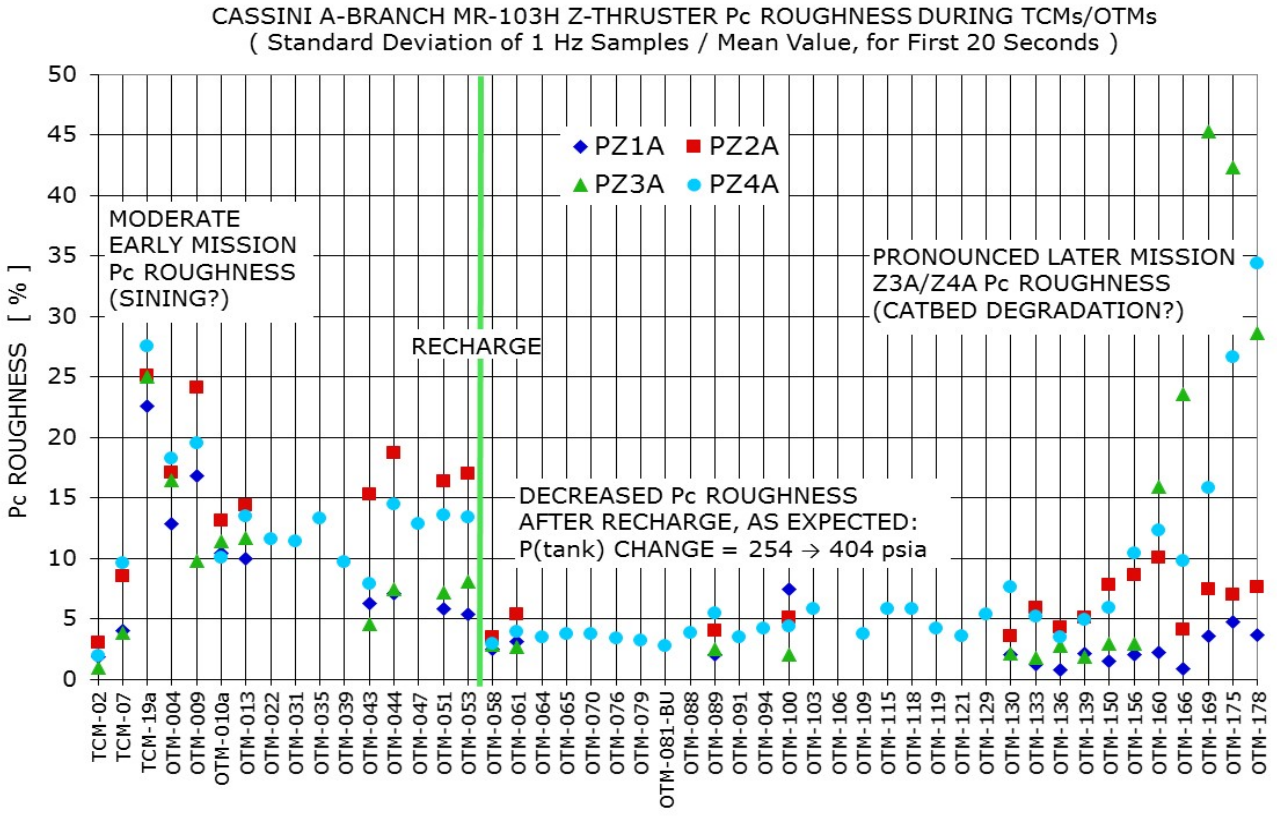

Figure 14: Cassini A-Branch Z-Thruster (MR-103H) Pc Roughness Proxy for all A-Branch RCS TCMs/OTMs

Figure 14 clearly shows early mission Pc roughness increasing from TCM-2 to TCM-7 to TCM-19a. Subsequent A-branch RCS OTMs displayed decreased roughness, perhaps another hint that sining fully explains early mission Pc roughness. Sining is also consistent with the marked decrease in roughness observed following the Monopropellant Tank Assembly (MTA) recharge on April 10, 2006, which fired PV40 to pressurize the MTA from the Recharge Tank Assembly (RTA). Specifically, recharge changed MTA pressure from 254 to 404 psia, a rather sizable increase in tank pressure. Sining is thought to be much less likely at higher tank pressures such as these. Following recharge, Pc roughness remained low until Z3A and Z4A began experiencing accelerating degradation in the second half of 2008. Given concerns about catbed void formation or perhaps even thruster burn-through, a permanent swap to B-branch RCS thrusters took place on March 12, 2009 ("as soon as feasible"). After a series of AACS checkouts, thrusters were deemed operational. They operated flawlessly for the next 8.5 years, closing out the highly successful Cassini mission.

Since root cause for A-branch degradation could not be determined (see Ref. 8), monitoring B-branch RCS thruster performance became paramount. Unfortunately, no Pc measurements were provided for back-up thrusters Z1B-Z4B, so only indirect techniques could be used to tease out any potential B-branch thruster degradation. Despite this limitation, the project performed admirably in monitoring potential issues with B-branch MR-103H thrusters. One generic finding from Ref. 8 was a recommendation to "share the pain" more equally between Z-thrusters and much less used Y-thrusters. To this end, so-called "Y-biases" were developed to manage reaction wheel life.[10] This technique would incur essentially zero Z-thruster firing, but it was presumed to come at a heavy hydrazine cost. Wonderfully, switching to Y-biases when possible actually saved ~30% in hydrazine usage vs. "normal" RWA biases, largely because multi-axis-use Z-thrusters were not "fighting each other" in setting optimal RWA wheel speeds. Given lean hydrazine margins during many years of tour, this change literally made this doubly extended mission viable, all while preserving Z1B-Z4B thruster life.

Many techniques were employed to monitor B-branch thruster health, particularly for Z1B-Z4B. Some of the health checks performed by AACS included observing thrust performance during RCS OTMs, looking for changes in off-pulsing duty cycle plots during RCS OTMs, and tracking actual impulse bits for Z1B-Z4B (and Y1B-Y4B) during RWA biases, using known spacecraft inertial properties, thruster on-times and cycles, and RWA wheel speeds. One further bold mitigation plan was developed by AACS as well—a substantial flight software change to enable spacecraft attitude control in so-called "mixed branch" mode, cherry-picking any given (healthy) A-branch or B-branch RCS thruster at will.[11] Even though mixed branch ops were never required on Cassini, thanks to excellent Z1B-Z4B and Y1B-Y4B performance through end of mission, this heroic effort by AACS should be lauded.

Propulsion engineers also found creative ways to monitor Z1B-Z4B health, despite not having any direct visibility into Pc roughness. Unlike A-branch (see Figure 8), B-branch line pressure telemetry was provided on Cassini. Presumably, if Pc spikes were large enough, upstream pressure fluctuations in line pressure telemetry or even in tank outlet pressure might be observable. As part of the campaign to monitor Z1B-Z4B health, Figure 15 was prepared—the now familiar standard deviation of 1-Hz samples divided by the average value over the first twenty seconds, this time for both B-branch line pressure and MTA outlet pressure, for all RCS TCMs/OTMs during the mission. Until the B-branch thruster swap in 2009, LV41 remained closed and only A-branch thrusters were used, so B-branch line pressure is not relevant before the swap. The two takeaways from Figure 15 are (1) Z1B-Z4B showed no signs of tank outlet pressure roughness or even line pressure roughness throughout the mission and, perhaps more interestingly, (2) Z3A and Z4A Pc roughness in OTM-169, OTM-175, & OTM-178 was increasingly visible, all the way upstream to the MTA outlet! This *post facto* observation suggests Z1B-Z4B thruster degradation could, indeed, have been readily spotted in B-branch line pressure telemetry if not MTA outlet pressure telemetry itself if it had occurred.

A summary of B-branch OTM ΔV performance is presented below in Figure 16. This is completely analogous to and on exactly the same scale as Figure 13. As with A-branch OTMs, ΔV errors were typically within ±3% for burns above 0.05 m/s, demonstrating excellent open-loop maneuver performance. Note the very large number of B-branch OTMs executed (cf. Figure 16 and Figure 13), a more symmetric error distribution around the zero line (due to improved deadband modeling), and increasing errors for smaller maneuvers, as expected. Tables 11-14 below summarize all 129 B-branch RCS OTMs between 2009 and 2017, as Table 10 did for A-branch OTMs. Though often exhausting and feeling like a marathon for sure (particularly given frequent evening, weekend, and holiday work, courtesy one Sir Isaac Newton), it was an honor and privilege guiding Cassini through 360 TCMs and OTMs over nearly two decades. The most complex planetary propulsion system ever flown earned its stripes at Saturn.

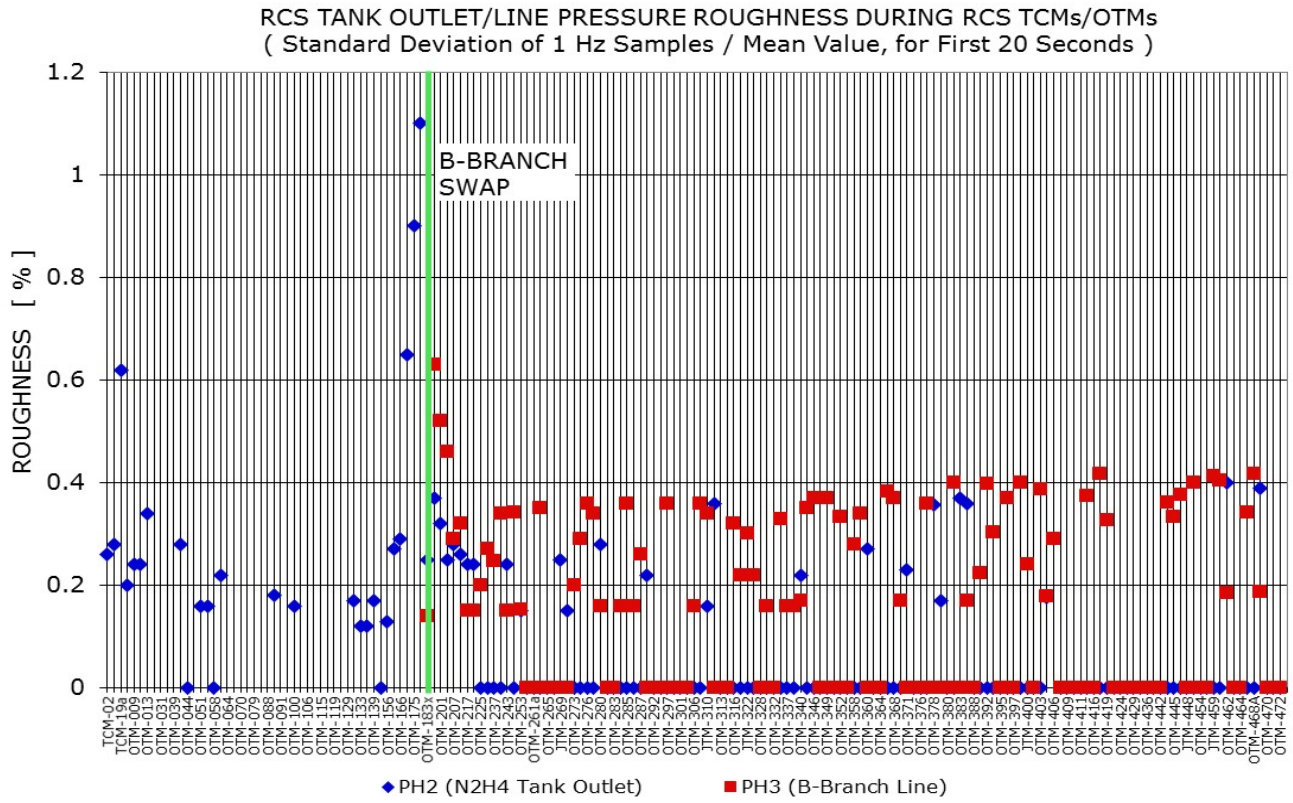

Figure 15: Cassini RCS Tank Outlet and B-Branch Line Pressure Roughness Proxy for all RCS TCMs/OTMs

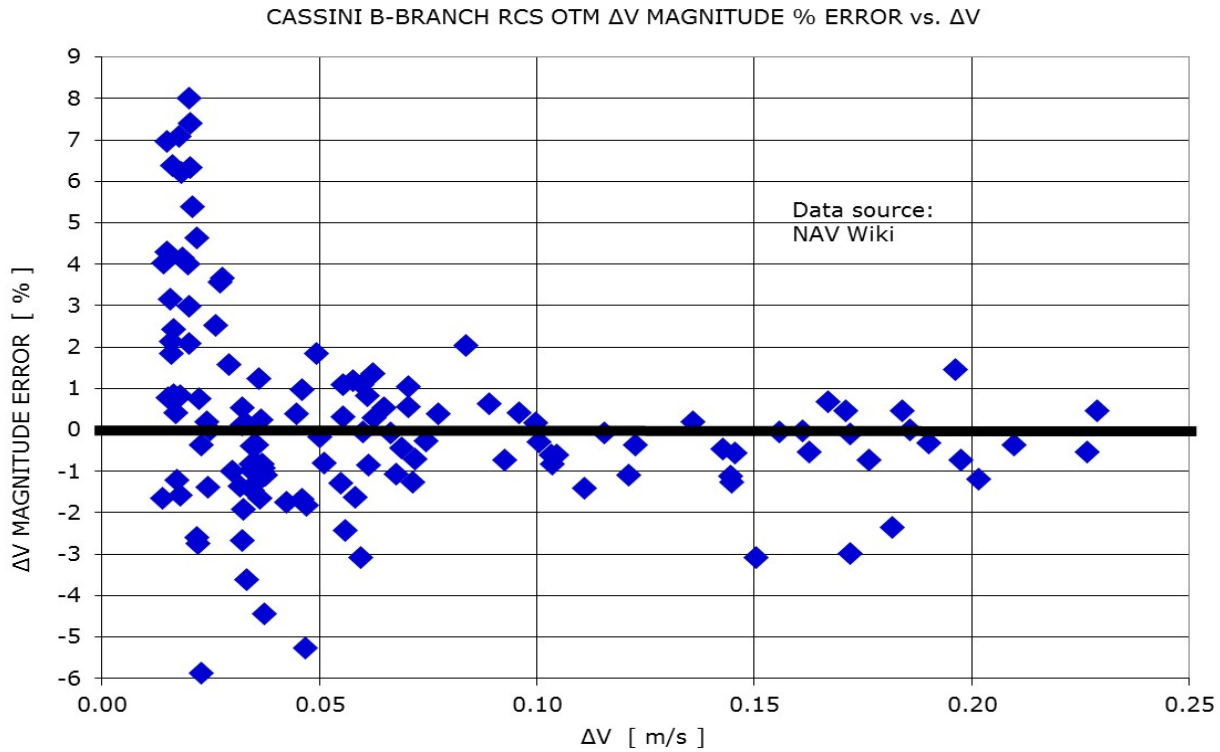

Figure 16: Cassini B-Branch RCS OTM ΔV Magnitude % Error vs. Executed ΔV

Table 11. Cassini RCS OTM Summary Table: OTM-183x thru OTM-313 (B-Branch RCS OTMs 3/18/09-3/24/12)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
183x	3/18/09	12.750	0.0218	7.41	T50-to-T51 apoapse burn; RCS B-Branch OTM checkout
196	5/18/09	37.625	0.0444	-5.27	Approach burn for T55 flyby on 5/21/09 @ 966 km altitude
201	6/15/09	22.000	0.0296	-1.00	T56-to-T57 apoapsis targeting burn for T57 flyby @ 955 km
204	7/1/09	9.125	0.0162	2.13	T57-to-T58 apoapsis targeting burn for T58 flyby @ 966 km
207	7/17/09	24.125	0.0318	-1.91	T58-to-T59 apoapsis targeting burn for T59 flyby @ 956 km
210	8/1/09	14.750	0.0226	0.76	T59-to-T60 apoapsis targeting burn for T60 flyby @ 971 km
217	10/9/09	130.125	0.1458	-3.09	Approach burn for T62 flyby on 10/12/09 @ 1300 km altitude
220	10/29/09	56.750	0.0669	-1.07	Approach burn for E7 flyby on 11/2/09 @ 99 km altitude
225	12/4/09	180.375	0.1992	-1.20	E8-to-T63 apoapsis targeting burn for T63 flyby @ 4848 km
232	1/9/10	28.375	0.0354	-1.23	Approach burn for T65 flyby on 1/12/10 @ 1074 km altitude
237	2/23/10	8.750	0.0160	6.97	T66-to-Rhea2 apoapsis targeting burn for R2 flyby @ 102 km
241	4/2/10	26.625	0.0339	-0.97	Approach burn for T67 flyby on 4/5/10 @ 7438 km altitude
243	4/18/10	38.250	0.0449	0.38	D2-to-E9 apoapsis targeting burn for E9 flyby @ 100 km
250	6/1/10	29.250	0.0365	-0.81	Approach burn for T69 flyby on 6/5/10 @ 2042 km altitude
253	6/18/10	18.875	0.0249	0.00	Approach burn for T70 flyby on 6/21/10 @ **878 km altitude**
256	7/4/10	14.625	0.0228	4.63	Approach burn for T71 flyby on 7/7/10 @ 1004 km altitude
261a	9/16/10	158.750	0.1750	-0.73	E11-to-T72 apoapsis targeting burn for T72 flyby @ 8178 km
264	10/15/10	165.125	0.1775	-2.35	T72-to-T73 periapse burn; 1st OTM **Cassini Solstice Mission**
265	11/8/10	157.500	0.1669	-2.99	Approach burn for T73 flyby on 11/11/10 @ 7926 km altitude
268	11/27/10	55.875	0.0651	0.52	Approach burn for E12 flyby on 11/30/10 @ 46 km altitude
269 b/u	12/1/10	146.875	0.1617	-0.54	E12 clean-up; b/u to save ΔV & avoid low RWA wheel speeds
270	12/8/10	9.500	0.0163	3.16	E12-to-E13 apoapsis targeting burn for E13 flyby @ 48 km
273	1/1/11	191.250	0.2091	-0.35	E13-to-Rhea3 apoapsis targeting burn for R3 flyby @ 70 km
274	1/8/11	27.875	0.0341	-0.82	Approach burn for R3 flyby on 1/11/11 @ 70 km altitude
276	2/1/11	13.125	0.0214	6.34	R3-to-T74 periapsis targeting burn for T74 flyby @ 3651 km
279	3/2/11	89.250	0.1002	-0.29	T74-to-T75 apoapsis targeting burn for T75 flyby @ 10053 km
280	4/15/11	13.125	0.0217	8.01	Approach burn for T75 flyby on 4/19/11 @ 10053 km altitude
281	4/22/11	35.500	0.0418	-1.75	Clean-up burn after T75 flyby targeting T76 flyby @ 1873 km
283	5/5/11	8.125	0.0147	4.02	Approach burn for T76 flyby on 5/8/11 @ 1873 km altitude
284	5/12/11	109.375	0.1198	-1.10	Clean-up burn after T76 flyby targeting T77 flyby @ 1359 km
285	5/24/11	30.000	0.0366	0.25	T76-to-T77 apoapsis targeting burn for T77 flyby @ 1359 km
286	6/17/11	8.750	0.0155	4.31	Approach burn for T77 flyby on 6/20/11 @ 1359 km altitude
287	6/24/11	133.625	0.1448	-0.57	Clean-up burn after T77 flyby targeting T78 flyby @ 5821 km
288	8/22/11	83.750	0.0919	-0.73	T77-to-T78 periapsis targeting burn for T78 flyby @ 5821 km
292	9/28/11	26.750	0.0331	0.17	Approach burn for E14 flyby on 10/1/11 @ 99 km altitude
294	10/5/11	66.125	0.0744	-0.26	Clean-up burn after E14 flyby targeting E15 flyby @ 1231 km
297	10/28/11	40.250	0.0464	0.97	E15-to-E16 apoapsis targeting burn for E16 flyby @ 496 km
300a	12/1/11	13.750	0.0218	5.39	E16-to-Dione3 apoapsis targeting burn for D3 flyby @ 99 km
301	12/9/11	11.250	0.0190	7.09	Approach burn for D3 flyby on 12/12/11 & T79 flyby 12/13/11
304	12/22/11	10.125	0.0168	2.44	T79-to-T80 apoapsis targeting burn for T80 flyby @ 29415 km
306	1/16/12	44.125	0.0501	1.85	T80-to-T81 apoapsis targeting burn for T81 flyby @ 31131 km
308	2/3/12	124.375	0.1362	0.18	Clean-up burn after T81 flyby targeting T82 flyby @ 3803 km
310 b/u	2/17/12	12.875	0.0205	4.00	T82 approach burn; b/u to avoid low RWA wheel speeds
312a	3/16/12	95.625	0.1040	-0.60	T82-to-E17 apoapsis targeting burn for E17 flyby @ 74 km
313	3/24/12	10.000	0.0173	6.39	Approach burn for E17 flyby on 3/27/12 @ 74 km altitude

Table 12. Cassini RCS OTM Summary Table: OTM-314 thru OTM-398 (B-Branch RCS OTMs 3/25/12-12/14/14)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
314	3/31/12	133.750	0.1429	-1.10	Clean-up burn after E17 flyby targeting E18 flyby @ 74 km
316	4/11/12	25.500	0.0314	-1.35	Approach burn for E18 flyby on 4/14/12 @ 74 km altitude
319	4/29/12	29.500	0.0348	-0.77	Approach burn for E19 flyby on 5/2/12 @ 74 km altitude
322 b/u	5/19/12	73.500	0.0775	0.39	T82 approach burn; b/u to avoid low RWA wheel speeds
325	6/3/12	31.750	0.0371	-1.09	Approach burn for T84 flyby on 6/7/12 @ 959 km altitude
328	7/21/12	161.250	0.1720	-0.09	Approach burn for T85 flyby on 7/24/12 @ 1012 km altitude
331	9/23/12	55.750	0.0607	-0.85	Approach burn for T86 flyby on 9/26/12 @ 956 km altitude
332	9/30/12	179.000	0.1895	-0.33	Clean-up burn after T86 flyby targeting T87 flyby @ 973 km
334	11/9/12	55.500	0.0600	-0.05	Approach burn for T87 flyby on 11/13/12 @ 973 km altitude
337	11/26/12	15.125	0.0214	-2.60	Approach burn for T88 flyby on 11/29/12 @ 1014 km altitude
338	12/2/12	22.500	0.0286	3.66	Clean-up burn after T88 flyby targeting T89 flyby @ 1978 km
340	2/13/13	27.375	0.0314	-2.67	Approach burn for T89 flyby on 2/17/13 @ 1978 km altitude
345	3/17/13	177.500	0.1859	0.00	R4-to-T90 apoapsis targeting burn for T90 flyby @ 1400 km
346	4/1/13	10.000	0.0167	0.84	Approach burn for T90 flyby on 4/5/13 @ 1400 km altitude
347	4/9/13	115.875	0.1223	-0.36	Clean-up burn after T90 flyby targeting T91 flyby @ 970 km
349	5/19/13	10.250	0.0170	0.41	Approach burn for T91 flyby on 5/23/13 @ 970 km altitude
350	5/27/13	44.875	0.0507	-0.80	Clean-up burn after T91 flyby targeting T92 flyby @ 964 km
352	7/7/13	53.750	0.0583	1.20	Approach burn for T92 flyby on 7/10/13 @ 964 km altitude
355	7/23/13	66.875	0.0716	-0.71	Approach burn for T93 flyby on 7/26/13 @ 1400 km altitude
358	9/9/13	30.250	0.0341	-1.45	Approach burn for T94 flyby on 9/12/13 @ 1400 km altitude
359	9/16/13	28.625	0.0320	-3.62	Clean-up burn after T94 flyby targeting T95 flyby @ 961 km
360	9/30/13	66.000	0.0705	-1.26	T94-to-T95 apoapsis targeting burn for T95 flyby @ 961 km
361	10/11/13	11.875	0.0193	4.15	Approach burn for T95 flyby on 10/14/13 @ 961 km altitude
364	11/28/13	7.750	0.0137	-1.64	Approach burn for T96 flyby on 12/1/13 @ 1400 km altitude
367	12/29/13	110.000	0.1153	-0.06	Approach burn for T97 flyby on 1/1/14 @ 1400 km altitude
368	1/5/14	99.375	0.1027	-0.61	Clean-up burn after T97 flyby targeting T98 flyby @ 1236 km
370	1/30/14	51.000	0.0545	-2.42	Approach burn for T98 flyby on 2/2/14 @ 1236 km altitude
371	2/5/14	85.625	0.0897	0.64	Clean-up burn after T98 flyby targeting T99 flyby @ 1500 km
373	3/3/14	19.750	0.0242	0.21	Approach burn for T99 flyby on 3/6/14 @ 1500 km altitude
376	4/4/14	50.500	0.0541	-1.29	Approach burn for T100 flyby on 4/7/14 @ 963 km altitude
377	4/11/14	33.375	0.0357	-4.45	Clean-up after T100 flyby targeting T101 flyby @ 2994 km
378	4/24/14	32.375	0.0356	-1.66	T100-to-T101 apoapsis target burn for T101 flyby @ 2994 km
379	5/14/14	17.250	0.0214	-5.88	Approach burn for T101 flyby on 5/17/14 @ 2994 km altitude
380	5/21/14	13.375	0.0204	2.10	Clean-up after T101 flyby targeting T102 flyby @ 3659 km
382	6/15/14	23.000	0.0282	3.57	Approach burn for T102 flyby on 6/18/14 @ 3659 km altitude
383	6/22/14	41.250	0.0451	-1.68	Clean-up after T102 flyby targeting T103 flyby @ 5103 km
385	7/17/14	29.000	0.0324	0.53	Approach burn for T103 flyby on 7/20/14 @ 5103 km altitude
388	8/18/14	29.000	0.0324	0.12	Approach burn for T104 flyby on 8/21/14 @ 964 km altitude
391	9/19/14	78.875	0.0854	2.04	Approach burn for T105 flyby on 9/22/14 @ 1400 km altitude
392	9/26/14	63.250	0.0664	-0.08	Clean-up after T105 flyby targeting T106 flyby @ 1013 km
394	10/21/14	32.125	0.0355	-0.36	Approach burn for T106 flyby on 10/24/14 @ 1013 km altitude
395	10/27/14	59.750	0.0631	1.35	Clean-up after T106 flyby targeting T107 flyby @ 980 km
396	11/22/14	200.625	0.1960	-0.73	T106-to-T107 apoapsis target burn for T107 flyby @ 980 km
397	12/7/14	33.750	0.0368	-0.92	Approach burn for T107 flyby on 12/10/14 @ 980 km altitude
398	12/14/14	164.375	0.1612	-0.02	Clean-up after T107 flyby targeting T108 flyby @ 970 km

Table 13. Cassini RCS OTM Summary Table: OTM-400 b/u thru OTM-472 (B-Branch RCS OTMs 12/15/15-9/15/17)

OTM	OTM Date [M/D/Y]	Burn Time [s]	ΔV [m/s]	ΔV Err. [%]	Primary Purpose of OTM
400 b/u	1/9/15	51.250	0.0561	1.08	T108 approach burn; b/u executed to save 0.11 m/s ΔV
401	1/14/15	231.500	0.2300	0.46	Clean-up after T108 flyby targeting T109 flyby @ 1200 km
403	2/9/15	25.375	0.0297	1.57	Approach burn for T109 flyby on 2/12/15 @ 1200 km altitude
405	3/4/15	100.250	0.0998	0.17	T109-to-T110 apoapsis target burn for T110 flyby @ 2274 km
406	3/13/15	17.625	0.0227	-0.35	Approach burn for T110 flyby on 3/16/15 @ 2274 km altitude
408	4/20/15	44.000	0.0462	-1.81	T110-to-T111 apoapsis target burn for T111 flyby @ 2722 km
409	5/4/15	11.875	0.0182	0.83	Approach burn for T111 flyby on 5/7/15 @ 2722 km altitude
410	5/11/15	60.250	0.0615	0.84	Clean-up after T111 flyby targeting Dione-4 flyby @ 516 km
411	6/8/15	58.125	0.0605	1.09	T111-to-D4 apoapsis targeting burn for D4 flyby @ 516 km
414	6/26/15	70.500	0.0708	0.57	D4-to-T112 apoapsis target burn for T112 flyby @ 10953 km
416	7/10/15	97.500	0.0963	0.41	Clean-up after T112 flyby targeting Dione-5 flyby @ 474 km
417	8/9/15	11.875	0.0178	-1.58	T112-to-D5 apoapsis targeting burn for D5 flyby @ 474 km
419	8/21/15	56.000	0.0572	-1.63	Clean-up after Dione-5 flyby targeting T113 flyby @ 1036 km
421	9/25/15	16.375	0.0216	-2.75	Approach burn for T113 flyby on 9/28/15 @ 1036 km altitude
424	10/11/15	32.500	0.0343	-0.38	Approach burn for E20 flyby on 10/14/15 @ 1839 km altitude
426	10/20/15	69.500	0.0711	1.05	E20-to-E21 apoapsis targeting burn for E21 flyby @ 49 km
429	11/5/15	112.125	0.1093	-1.40	E21-to-T114 apoapsis target burn for T114 flyby @ 11920 km
431	11/16/15	105.250	0.1028	-0.83	Clean-up after T114 flyby targeting E22 flyby @ 4999 km
436	1/12/16	34.375	0.0366	1.24	Approach burn for T115 flyby on 1/16/16 @ 3817 km altitude
439 b/u	1/29/16	9.875	0.0161	1.843	Approach burn for T116; b/u executed due to DSN problems
442	2/13/16	9.375	0.0154	0.79	Approach burn for T117 flyby on 2/16/16 @ 1018 km altitude
443 b/u	2/20/16	68.500	0.0685	-0.44	T117 clean-up; b/u to lower nav errs for Enceladus occultation
445	4/1/16	62.250	0.0627	0.30	Approach burn for T118 flyby on 4/4/16 @ 990 km altitude
446	4/7/16	171.750	0.1680	0.68	Clean-up after T118 flyby targeting T119 flyby @ 971 km
448 b/u	5/4/16	11.250	0.0170	-1.21	Approach burn for T119; b/u to save ΔV & reaction wheel life
450	5/22/16	23.125	0.0268	2.53	T119-to-T120 apoapsis target burn for T120 flyby @ 975 km
454	7/22/16	47.875	0.0500	-0.18	Approach burn for T121 flyby on 7/25/16 @ 976 km altitude
455	7/28/16	194.750	0.1848	0.46	Clean-up after T121 flyby targeting T122 flyby @ 1599 km
459 b/u	8/20/16	53.125	0.0555	0.32	T122-to-T123 periapsis; b/u to save ΔV & reaction wheel life
460	9/23/16	21.500	0.0241	-1.39	Approach burn for T123 flyby on 9/27/16 @ 1736 km altitude
462	10/5/16	181.250	0.1718	0.47	T123-to-T124 periapsis target burn for T124 flyby @ 1582 km
463	11/10/16	12.500	0.0195	6.22	Approach burn for T124 flyby on 11/14/16 @ 1582 km altitude
464	11/17/16	151.000	0.1420	-0.47	Clean-up after T124 flyby targeting T125 flyby @ 3223 km
468	12/24/16	239.250	0.2253	-0.54	T125-to-T126 apoapsis target burn for T126 flyby @ 979 km
468a	2/22/17	207.375	0.1990	1.46	T125-to-T126 periapsis target burn for T126 flyby @ 979 km
469	4/18/17	58.000	0.0577	-3.07	Approach burn for T126 flyby on 4/22/17 @ 979 km altitude
470	4/24/17	165.250	0.1555	-0.06	Clean-up errors after T126 flyby; **first proximal orbit**
471	5/10/17	14.250	0.0207	2.98	Proximal orbit targeting maneuver #2 of 3 (OTM-470 = #1)
472	7/15/17	153.125	0.1429	-1.25	Proximal orbit targeting maneuver #3 of 3

VI. Pressure Transducer Drift

Many JPL missions (Voyager, TOPEX-Poseidon, and Galileo, to name a few) have experienced linear pressure transducer drift during the course of their multi-year missions. Frustratingly, this phenomenon has befuddled attempts to understand pressurization system behavior, propellant consumption, and maneuver performance during mission operations. Therefore, great pains were taken on Cassini to provide drift-free pressure transducers. This included modifications to the electronic circuitry, since this was typically implicated as the culprit for prior sensor drift. Specifically, an operational amplifier (op amp) in the pressure transducer supply electronics was found to drift linearly vs. time for some prior JPL missions. This single component instability is sufficient to explain the pressure transducer drift seen in flight on Voyager, TOPEX-Poseidon, and Galileo.

Unfortunately, in the spaceflight environment, there are often no independent reference points to assess actual pressure, so it is impossible to determine which transducers are drifting, even for multiple sensors measuring the same physical quantity. There was one happy exception during the Cassini mission, however. When PV40 was fired open on April 10, 2006, transferring helium mass from the Recharge Tank Assembly (RTA) to the Monopropellant Tank Assembly (MTA), hydrazine tank pressure increased from 254 to 404 psia. This change in pressure was exactly as predicted, to the DN (Data Number) discretization level of 1.88 psia/DN. If Cassini's MTA pressure sensor were drifting at rates similar to those observed on prior JPL missions, a discrepancy of 4-5 DN would have been observed at recharge, so precise DN agreement at least qualitatively suggests much lower sensor drift rates on Cassini.

Pressure transducer drifts on previous missions were discovered in flight by differencing the output of two, independent sensors that measured the same pressure and seeing how the difference between the two measurements grew over time. Therefore, spacecraft that had no redundant pressure measurements (such as Deep Space One) were not useful for assessing drift. The maximum drift rate observed in flight was quite consistent among Galileo, Voyager, and TOPEX-Poseidon, roughly 0.24%-0.32% of full scale per year. This is consistent with the op amp anomalous performance mentioned above. The specification for maximum allowable pressure transducer drift is typically 2.0% of full scale over the life of the mission, so this means prior JPL missions experienced out-of-spec drift after 6-8 years. Fortunately, Cassini drift was miniscule to virtually non-existent, even after two decades, as may be verified below.

Figure 17 represents the difference of two Cassini NTO tank pressure measurements, PO2 and PO1, as a function of mission time between launch and the end of mission. Note there was an offset present at launch; in fact, this offset was well known before launch. PO2 was known to be more representative of actual pressure values, so the measurement from PO1 was not used in flight, except in the assessment of pressure sensor drift. The difference between PO2 and PO1 was nearly flat vs. time, suggesting the Cassini NTO tank pressure transducers did not drift much at all, even over two decades. The slope in Figure 17 is just +0.0080% of full scale per year, 30-40 times smaller than the inferred drift rates on Galileo, Voyager, and TOPEX-Poseidon. Despite this low value, the accumulated effect after twenty years was just above the 1 DN threshold, the discretization limit for analog-to-digital telemetry. It is possible pressure transducer was present, but at this minute level, it had no impact on mission operations.

Figure 18 is the MMH-tank analogue to Figure 17. Note the offset between PF2 and PF1 was much smaller than it was for the NTO tank, typically 1±1 psid. Even better, the fuel pressure transducers apparently did not drift with respect to each other at all. Naturally, if PO1/PO2 or PF1/PF2 drifted at the same rate with respect to one another, this would be unobservable in Figures 17 and 18, since only relative drift is discernible using this technique. The slope in Figure 18 is a mere –0.00073% of full scale per year, 330-440 times smaller than the inferred drift rates on prior spacecraft! As listed in Table 2, Cassini pressure transducers were provided by Gulton-Statham, with flight heritage from ACTS (Advanced Communications Technology Satellite) and the DMSP (Defense Meteorological Satellite Program). They were requalified following design modifications, and their excellent performance in flight greatly aided propulsion mission operations for Cassini.

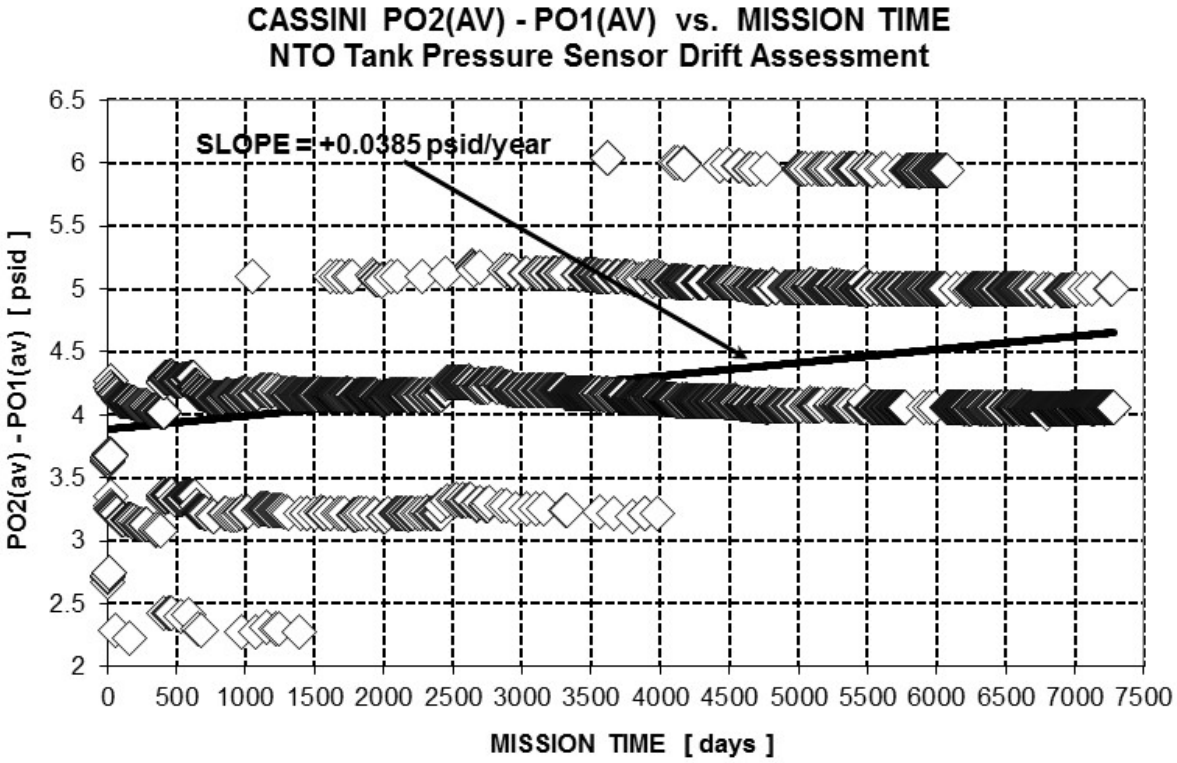

Figure 17: Cassini NTO Tank Pressure Relative Transducer Drift vs. Mission Time

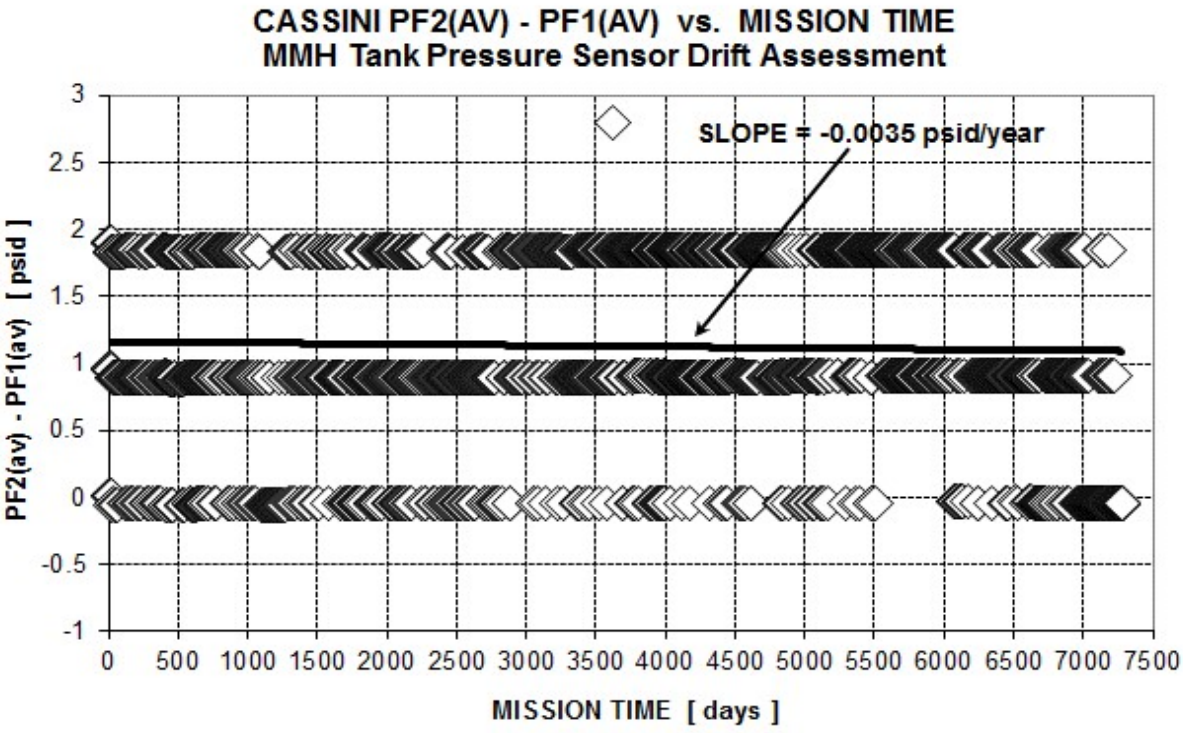

Figure 18: Cassini MMH Tank Pressure Relative Transducer Drift vs. Mission Time

Another way to investigate pressure transducer drift was developed anew for Cassini. If one were to linearly fit the data of Figures 17-18, stopping at an arbitrary point in the mission, this would be the inferred relative drift rate up to that point in the mission. In fact, this value may be calculated for each of the nearly 3000 points in Figure 17 or Figure 18 and then plotted vs. mission time (see Figure 19). Essentially, this shows the evolution vs. time of the inferred drift rate. For classic pressure transducer drift (i.e., Voyager, Galileo, or TOPEX-Poseidon), the drift rate settled fairly quickly to a constant, non-zero value, after initial transients averaged out. In contrast, Cassini inferred drift rates featured multiple "zero-drift" crossings (at least for PF1/PF2), generally decreased vs. mission time, and did not approach a non-zero asymptote, even after twenty years of operations. This type of "meta" plot may be a more useful way to discern pressure transducer drift on long missions, but the bottom line is Cassini had apparently nearly drift-free pressure transducers, despite the lengthy duration of its spaceflight.

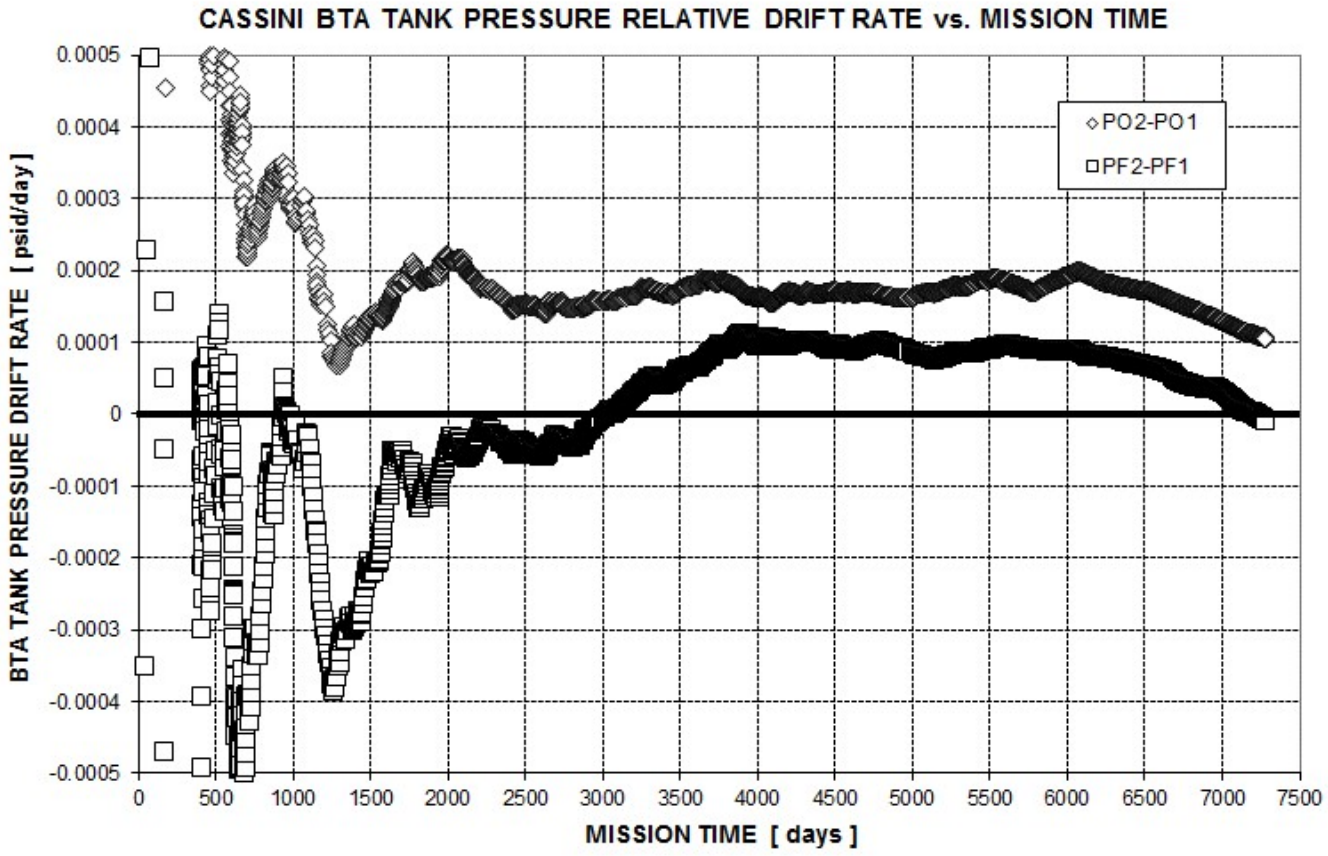

Figure 19: Cassini NTO & MMH Tank Pressure Relative Transducer Drift Rate vs. Mission Time

One final assessment of potential pressure sensor drift on Cassini was possible, this time by differencing N_2H_4 tank pressure (PH2) and B-branch line pressure (PH3). However, until LV41 was opened on March 10, 2009 in preparation for the B-branch thruster swap, these two transducers were not measuring the same physical quantity. Still, this means 8.5 years of relative pressure transducer drift between PH2 and PH3 could be investigated from 2009 through the end of mission. Figure 20 displays the quantity [PH2 – PH3] as a function of mission time following the B-branch swap. A linear fit of the data shows changes far below the ±1 DN range spanned by the data points in the plot. This again strongly suggests no relative drift between PH2 and PH3 for 8.5 years if not twenty years. Kudos are due the CPMS development team for delivering exemplary pressure transducers for this multiply extended mission.

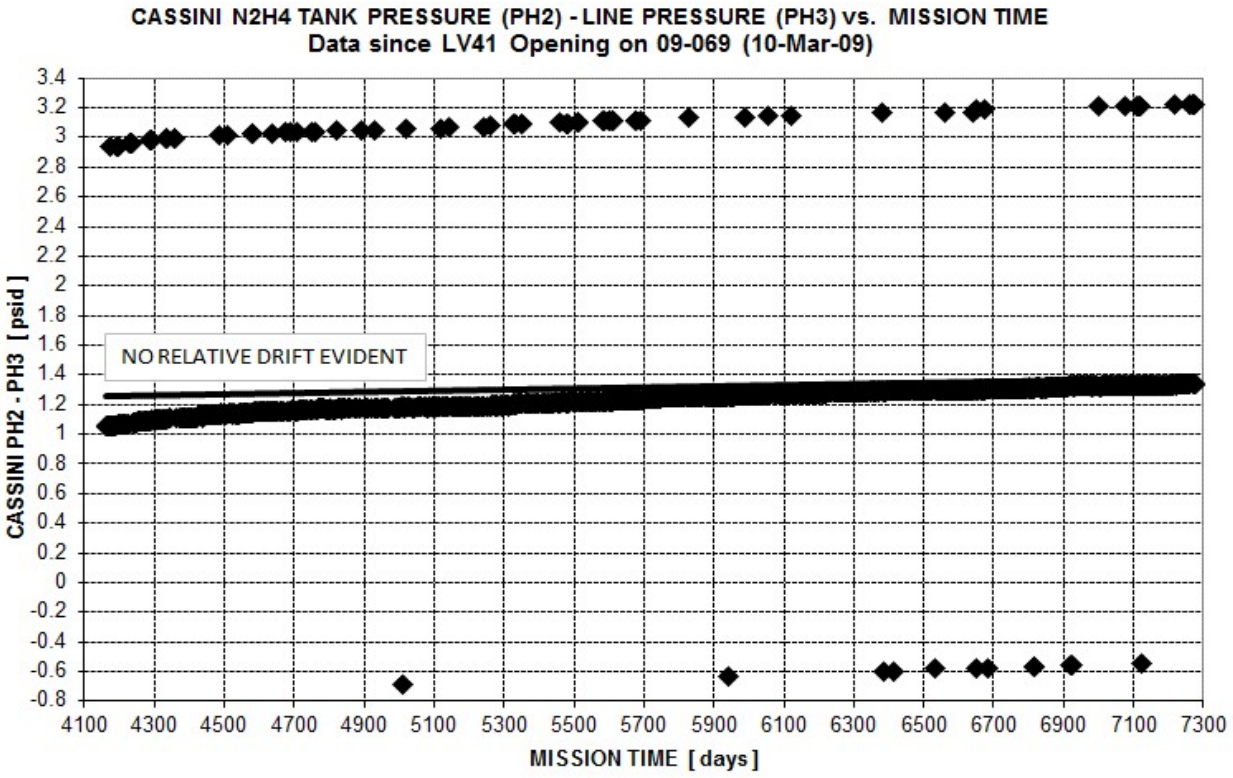

Figure 20: Cassini N_2H_4 Tank/B-Branch Line Pressure Relative Transducer Drift Rate vs. Mission Time

VII. Latch Valve Back-Pressure Relief Performance

To be more fault tolerant with respect to external propellant leakage, it is often good practice to have multiple seals protecting against this failure mode. As such, Cassini flew nearly twenty years with downstream NTO and MMH latch valves (LV21 and LV31, respectively) closed, except during main-engine burns using REA-A. By design, however, this created a "liquid lock" condition between LV21/LV31 and the REA-A engine valves, one which had to be accommodated in order to avoid potential line rupture, even with modest increases in propellant line temperature. This was achieved by selecting latch valves with back-pressure relief (BPR) functionality. Between launch in 1997 and the RCS thruster branch swap in 2009, LV41 remained closed as well, another liquid-lock condition for Cassini over a dozen years. Like LV21/LV31, BPR functionality was included for LV41 (and LV40) as well. Since LV40 only closed briefly during one spacecraft safing event, and moreover there was no pressure telemetry provided downstream of LV40 (see Figure 8), LV40 BPR was never observed during the Cassini mission.

The LV41 BPR opening or "cracking" spec was 100 ± 20 psid, and in-flight BPR for this latch valve occurred in flight a mere seven times and always within spec—at least when observable. The first three LV41 BPR events transpired during the initial four months of the mission, caused by post-launch heating, TCM-1, and TCM-2. However, this early in the mission, tank pressures were 368-372 psia, so spec BPR of LV41 would not be expected until downstream pressures reached 448-492 psia. Unfortunately, downstream line pressure pegged at 464 psia (i.e., 255 DN for these 8-bit measurements) for all three of these BPR events, so the true BPR point was not known. Later LV41 BPR events suggest spec violations likely did not occur during these first three BPR events, but there is no way to know for sure. For this reason, the first three points in Figure 21 below are denoted with arrows and question marks. Fortunately, the four BPR events that followed in succession (caused by a Venus-2 flyby turn, TCM-9, a spacecraft turn, and the RCS B-branch catbed heater test, respectively) all took place with lower tank pressures, non-pegged line pressures, and in-spec BPR performance, as may be verified in Figure 21.

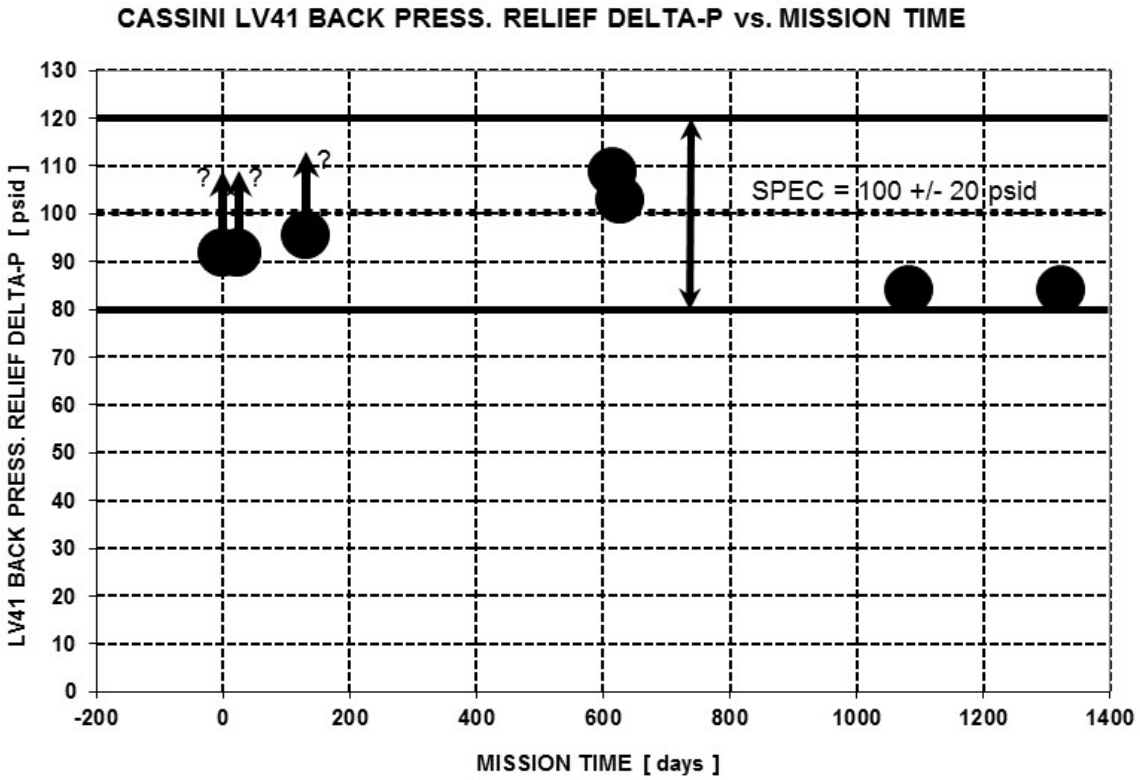

Figure 21: Cassini LV41 Back-Pressure Relief ΔP vs. Mission Time

Both LV21 and LV31 have a wider BPR spec than LV40 or LV41, 90 +40/-30 psid. During Cassini's two decades in flight, LV21 experienced BPR about 380 times, within a very narrow range of 108-120 psid, and there were no time trends in the pressure differential for BPR. In contrast, LV31 underwent BPR about 170 times, over a much broader range of 50-90 psid, and out of spec 43% of the time. Fortunately, this out-of-spec performance was in the safe direction, towards easier BPR. As such, it had no consequences for mission operations or risk to the spacecraft. Figure 22 below is the LV21/LV31 equivalent to Figure 21, representing nearly two decades of excellent LV21/LV31 BPR performance. Incidentally, LV40 was left open for the entire mission, even after the B-branch swap. Given non-leaking A-branch RCS thruster valves and unobservable LV40 BPR, this was deemed to be lower risk for the mission.

One useful exercise is to compare in-flight BPR performance with ground test performance. There is no spec for BPR closing or reseat Δp, only opening or cracking Δp. Moreover, in flight there was no way to accurately determine reseat Δp. Despite this limitation, reseat Δp was tabulated during ground test for LV21, LV31, and LV41, so these ranges are included in Table 14 below for completeness. The most interesting observation from Table 14 is LV21 BPR cracking Δp was low and more variable on the ground vs. flight, the exact opposite of LV31 BPR cracking Δp.

Table 14: Comparison of LV21 & LV31 & LV41 Ground-Test vs. In-Flight Performance

Latch Valve	Flight Relief ΔP Range [psid]	Ground Relief ΔP Range [psid]	Spec. [psid]	Flight Reseat ΔP Range [psid]	Ground Reseat ΔP Range [psid]	Spec. [psid]
LV41	84 to 108+ (?)	100 to 120	100 ± 20	---------	75 to 105	N/A
LV21	108 to 120	70 to 110	90 +40/-30	---------	95 to 105	N/A
LV31	50 to 90	90 to 98	90 +40/-30	---------	85 to 90	N/A

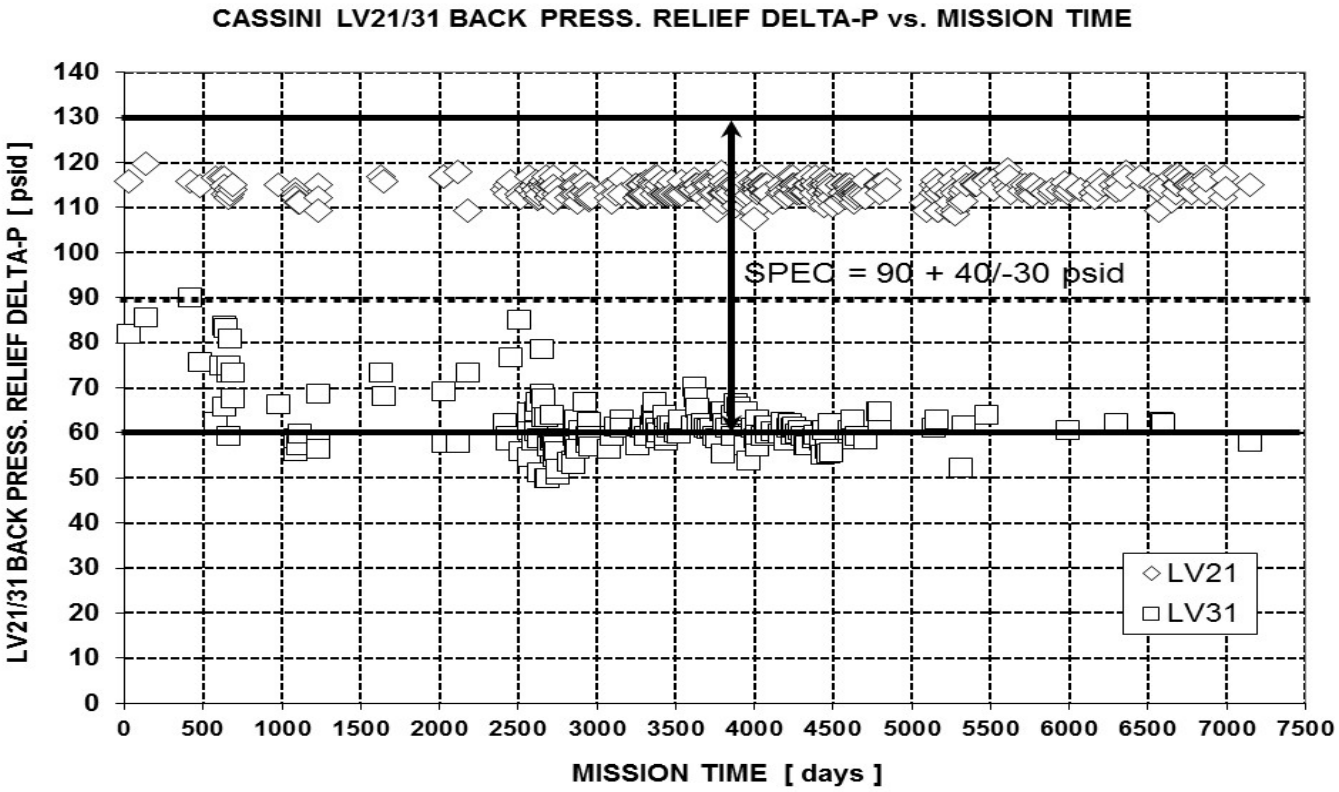

Figure 22: Cassini LV21/LV31 Back-Pressure Relief ΔP vs. Mission Time

Finally, it was noticed in flight that LV21 and LV31 BPR occurred at a higher Δp during periods of rapid line pressure rise. This is as expected, given BPR is not an instantaneous phenomenon. One way to investigate the effect of line pressure rise rate on BPR, at least qualitatively, is to "bin" similar BPR events and sort them by increasing line pressure rise rate. In-flight activities which caused line pressures to rise, in order of increasing rate, were (1) closing the main-engine cover, (2) spacecraft turns, (3) RCS OTMs/TCMs, (4) OTM/TCM tracking passes with no maneuver executed (but which still included line heating due to REA-A and REA-B heater turn-on), (5) main-engine OTMs/TCMs, and (6) the Propellant Gauging Test (PGT), a 2012 exercise designed to heat up the CPMS as quickly as possible before rapidly cooling the tanks in an effort to determine remaining NTO/MMH masses via thermal capacitance modeling. Figures 23 and 24 below display BPR cracking Δp for LV21 and LV31, respectively, as a function of these increasing line pressure rise rate bins. As expected, there does appear to be good correlation between increasing BPR cracking Δp and increasing line pressure rise rate, for both LV21 and LV31. Incidentally, a maneuver block change early in 2012 precluded BPR during future back-up OTM windows (bin #4 above). This was accomplished by setting a flag to "undo" REA-A and REA-B heater changes after five minutes if the prime OTM had already executed. This change was requested by propulsion as a simple way to reduce the number of BPR events.

In summary, Cassini latch valve BPR for LV41, LV21, and LV31 was nominal over nearly twenty years of mission operations. LV40 BPR was never tested in flight, since this latch valve remained open throughout the mission, except for one brief closure associated with a spacecraft safing event. These four latch valves were the only latch valves with trapped propellant liquid downstream, since LV22 and LV32, feeding REA-B, were never opened to prime the back-up main engine (normally closed PV2Z and PV32 would have to have been fired to prime REA-B as well, as may be verified in Figure 8). Though BPR anomalies were certainly unlikely or perhaps even non-credible on Cassini, it was comforting to see BPR functionality work so well in flight.

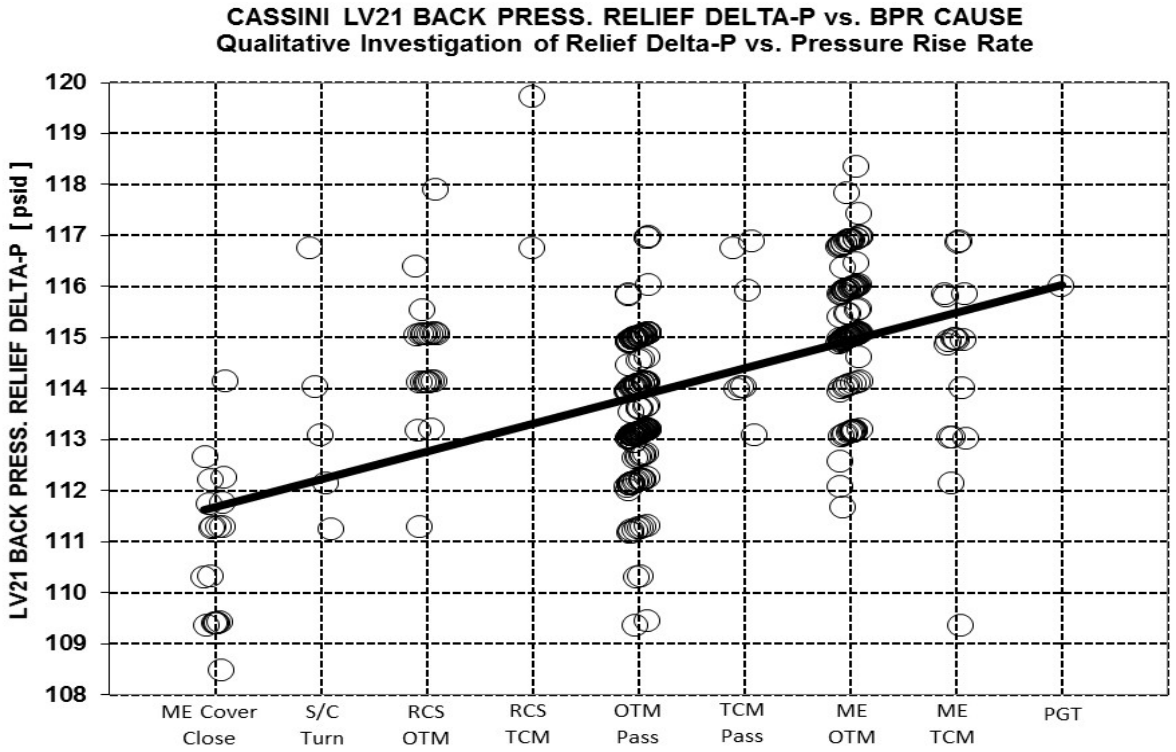

Figure 23: Cassini LV21 Back-Pressure Relief ΔP vs. Pressure Relief Cause

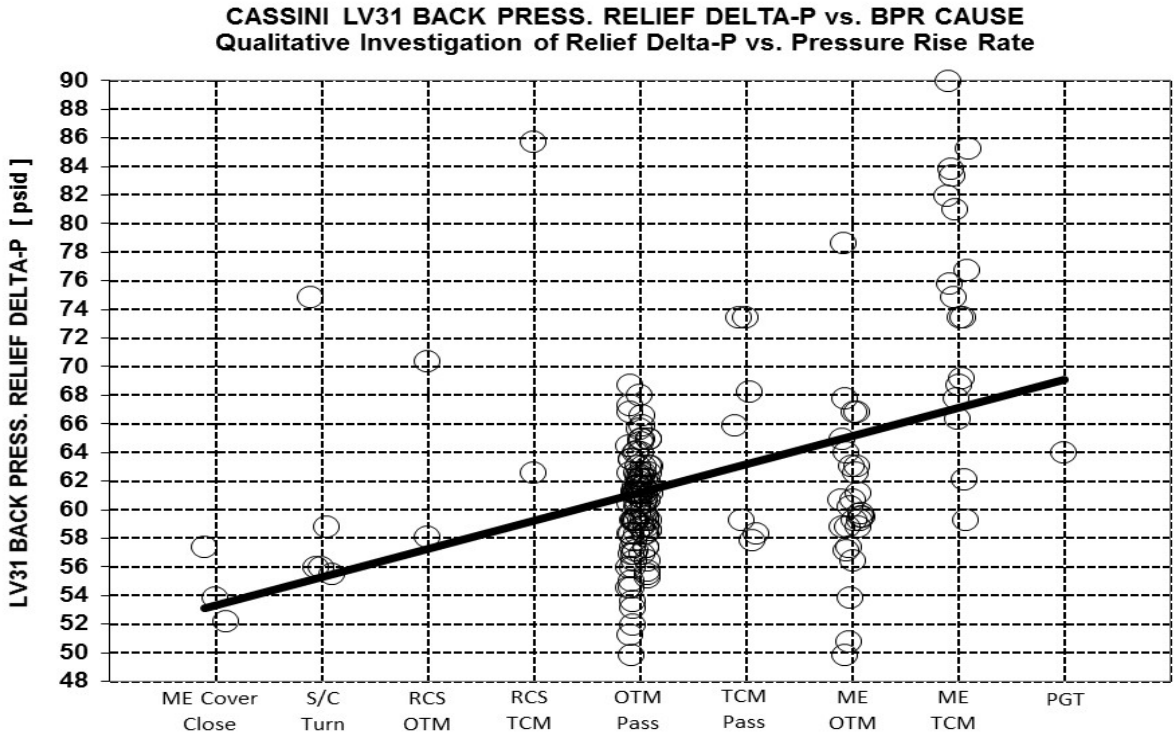

Figure 24: Cassini LV31 Back-Pressure Relief ΔP vs. Pressure Relief Cause

VIII. Monopropellant System Hydrazine & Helium Budgets

An initial accounting of the hydrazine mass remaining between launch and end of mission is displayed below in Figure 25. There are two methods that can be used to estimate spacecraft hydrazine mass. First, a hydrazine consumption model was used to attempt to "bean-count" every drop of hydrazine leaving the spacecraft. This model was of limited use due to the inaccuracies in the pulse-mode consumption model for the MR-103 thruster, but over short time scales, it was the only way to effectively estimate hydrazine usage for a given propulsive event. For every JPL mission flying this class of thruster, hydrazine usage has been overpredicted using consumption modeling, which is in the conservative direction with respect to propellant depletion. An alternative method is to simply calculate the hydrazine mass remaining from telemetered values of tank pressure and temperature using a thermodynamic model. This method should offer the greatest accuracy over the long term, as long as pressure transducers are not drifting (as was verified above). The data of Figure 25 were generated using the original pre-launch MR-103 consumption or "throughput" model, along with the thermodynamic or "tank" model. Nearly 3000 hydrazine-consuming events were tabulated over the twenty-year Cassini mission, leading to the fine detail evident in the data of Figure 25.

As anticipated, the consumption model consistently overpredicted N_2H_4 usage vs. the tank model, over the entire mission. By the time of Cassini's plunge into Saturn, the original throughput model suggested 132 – 18 = 114 kg of N_2H_4 was expended, or 114/132 = 86% of the total launch load including unusable hydrazine (more on unusable monopropellant and bipropellant below). However, the more believable tank model indicated only 132 – 36 = 96 kg of N_2H_4 was used, or 96/132 = 73% of the launch load. Equivalently, it appears the original consumption model overpredicted hydrazine usage by (114/96 – 1) = 19% vs. the tank model. This discrepancy is actually very typical; some JPL missions had pulse-mode consumption errors above 30% vs. tank models (e.g., Deep Space One). Most of Cassini's MR-103H pulses were 125 ms in duration, but during VVEJGA cruise, A-branch thrusters fired many thousands of pulses with quite short on-times between 7 and 77 ms. Consumption modeling errors for very small pulses can be rather sizable, so 19% agreement between the two curves in Figure 25 is actually not too bad.

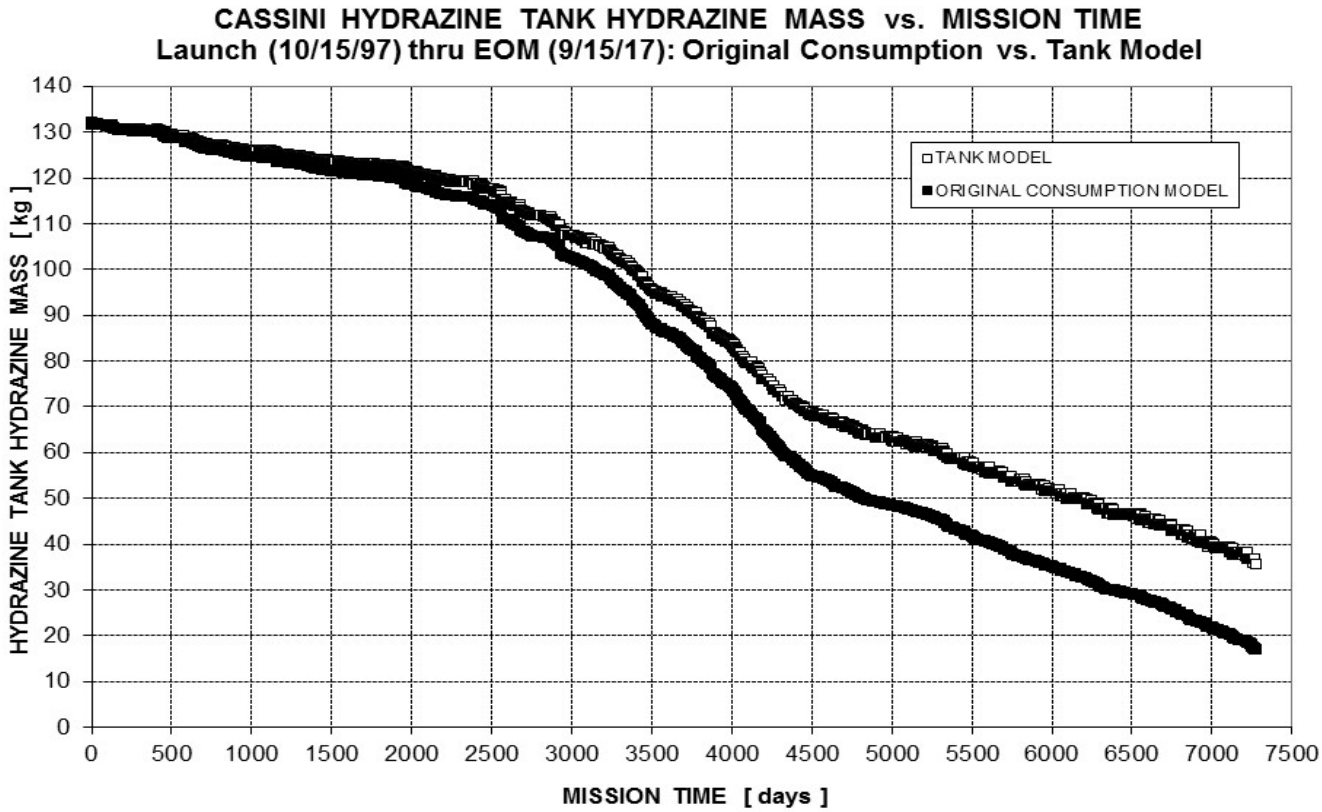

Figure 25: Cassini MTA Hydrazine Mass vs. Mission Time (Original Consumption vs. Tank Models)

It is clear from Figure 25 that Cassini hydrazine usage was not linear vs. mission time. Specifically, cruise used hydrazine at a very low rate (mission days 0-2451), while the four-year prime mission was much more propulsively intensive (mission days 2451-3940), as expected. The two-year equinox tour (mission days 3940-4732) used hydrazine at very nearly the same rate as the prime tour, at least initially. In contrast, the seven-year solstice tour (mission days 3942-7275) expended N_2H_4 at a much lower rate. Putting it bluntly, it is evident from Figure 25 that prime and equinox consumption rates were not sustainable for the solstice mission, including Cassini's Grand Finale. Thankfully, a number of hydrazine-saving initiatives were undertaken for solstice, such as the Y-biasing technique mentioned above, along with reduced RWA bias frequency and minimization of zero rpm RWA crossings during RBOT (RWA Bias Optimization Tool) design by AACS.[12] Earlier hydrazine conservation techniques continued to be employed as well, including performing RCS TCM/OTM yaw and roll turns under RWA control, executing main-engine roll turns on RWAs (spacecraft power margin was insufficient to perform main-engine yaw turns on RWAs), and using low-rate RCS turns when possible, especially for main-engine yaw turns. The solstice mission also was not as scientifically intensive as the prime mission or equinox; this saved copious N_2H_4 as well.

Given the disagreement between the two curves in Figure 25, a better consumption model was developed early during the Cassini mission, thanks to the extensive efforts of two Cassini summer students.[6] In-flight estimates of MR-103H impulse bit (IBIT) were generated for a whopping 1764 pulse-mode firings of thruster couples for spacecraft pitch, yaw, and roll control. On-times for these pulses were 7-77 ms, inclusive, and numerous inferred in-flight IBIT values for 125-ms pulses "anchored" the curve fit further in Ref. 6. For convenience, a summary plot for the entirety of this data set from Ref. 6 is reproduced below as Figure 26. The key point from all this work is the Wright valve ground test model used to generate the throughput curve in Figure 25 is conservative with respect to in-flight performance for spacecraft using Moog valves, such as Cassini and Voyager. Therefore, the red curve fit below was substituted into the original consumption model to more accurately track hydrazine consumption through Moog valve MR-103 thrusters. Moog valve results during ground testing are quite similar to red curve below, as expected.

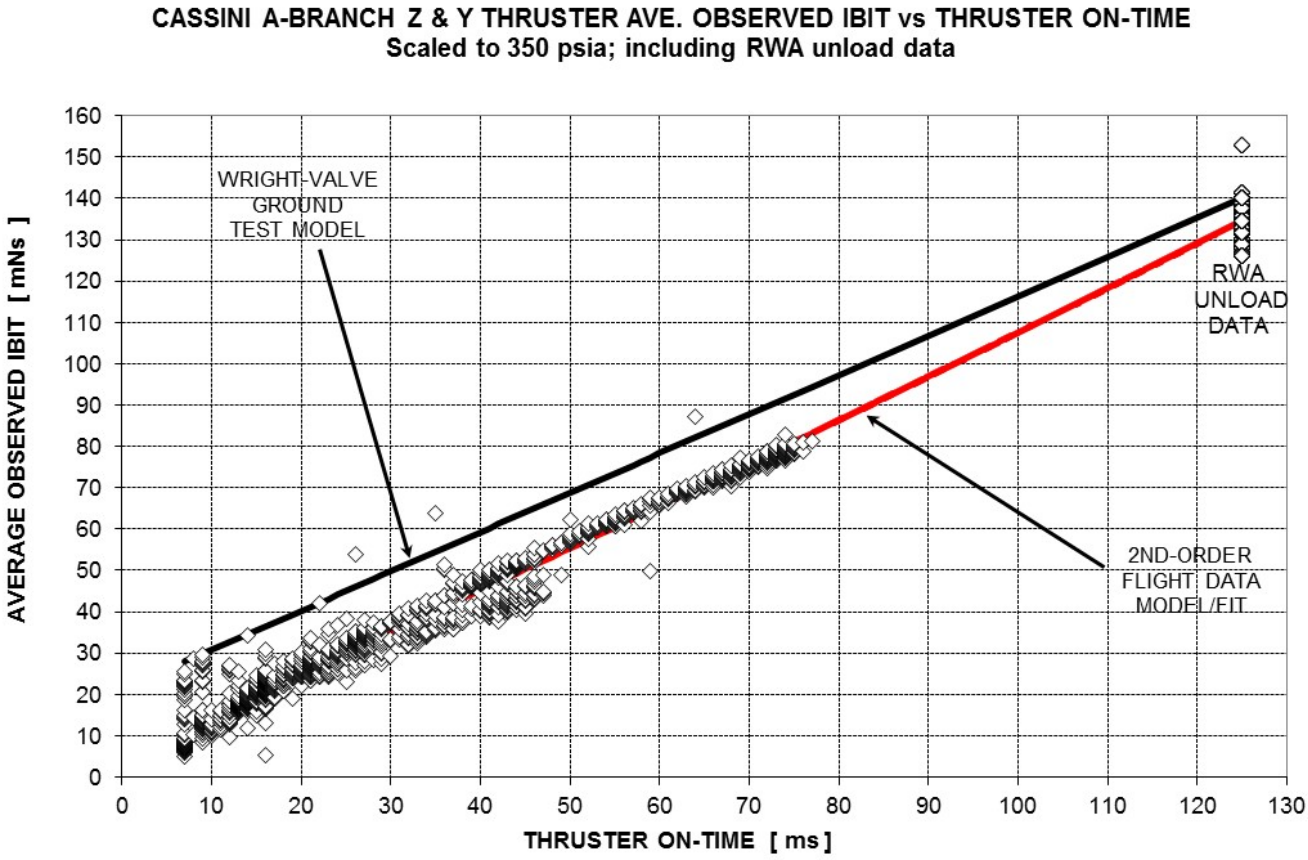

Figure 26: Cassini A-Branch RCS Thruster Observed Impulse Bit in Pulse Mode vs. Thruster On-Time

When substituting this improved IBIT vs. thruster on-time model (based on in-flight Moog valve data) into the relevant portion of the original Cassini throughput model, the discrepancy between tank and consumption models decreased, as expected. Quantitatively, this IBIT-modified consumption model suggested 132 – 21 = 111 kg of N_2H_4 was expended through end of mission, or 111/132 = 84% of the total launch load including unusable hydrazine. Equivalently, it appears this consumption model overpredicted hydrazine usage by (111/96 – 1) = 16% vs. the tank model. Admittedly, this is a modest improvement from the original consumption model error of 19%, but this is as expected, since most Cassini hydrazine was used in RCS OTMs (i.e., nearly steady-state operation) and for RWA biases, Titan flybys, and spacecraft turns, all of which fired 125-ms pulses. Figure 27 below shows the slight improvement in the discrepancy between tank and throughput models with this IBIT modification (cf. Figure 27 and Figure 25).

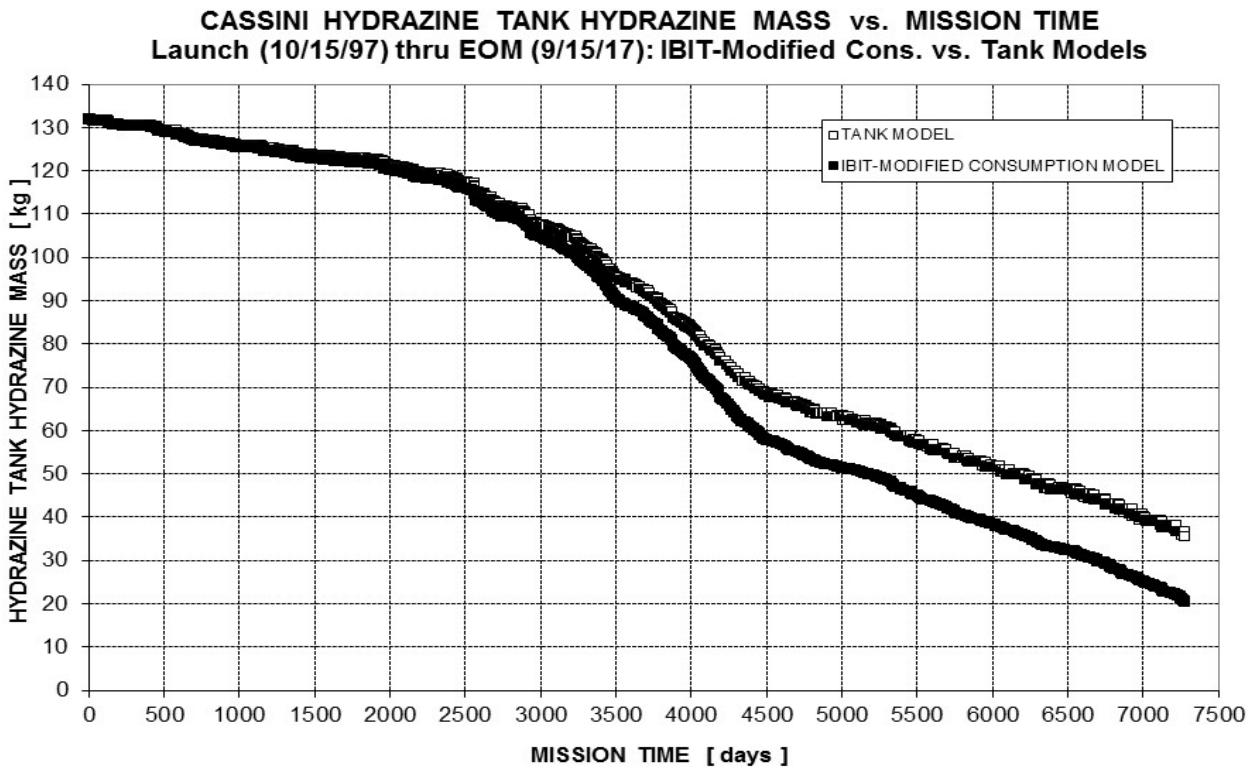

Figure 27: Cassini Hydrazine Tank Hydrazine Mass vs. Mission Time (w/ Improved IBIT Model for Moog Valves)

Even with the IBIT modification presented in Figure 27, the disagreement between tank and consumption models was still unacceptably large for hydrazine accounting purposes, especially as hydrazine margins became tight. One final attempt was made to align these two hydrazine mass models, way back in the summer of 2006. The hydrazine consumption equation used to generate the bottom curve in Figure 27 has very few "knobs" which may be tweaked. Specifically, other than a well-known linear dependence on tank pressure and the IBIT vs. on-time curve fit itself, the only other input affecting propellant consumption is specific impulse. Pulse-mode Isp is notoriously difficult to model, however, so minor tweaks to Isp to force the bean-counting model to agree with the tank model seemed quite justifiable. Prior missions modified pulse-mode Isp in consumption models as well, sometimes without a clear paper trail, unfortunately. Figures 25 and 27 were included above to capture Cassini N_2H_4 consumption modeling history.

Cassini fired hundreds of thousands of MR-103H pulses over the years, so calculating propellant consumption for each pulse was prohibitively laborious. Rather, each separate propulsive event (RCS OTM, RWA bias, spacecraft turn, Titan flyby, etc.) was assessed as a whole, using crude estimates of average Isp for various activities. Steady-state Isp was modeled at 217 s, while pulse-mode operations were "binned" with rough Isp values based on ground test data. Representative Isp values initially used were 120 s for limit duty cycling (less than 1% duty

cycles), 140 s for RWA biases (1-10% duty cycles), and 180 s for spacecraft turns on RCS and Titan flybys (10-80% duty cycles). Determining the best way to fine-tune these Isp values to force the consumption and tank models to agree was a black art at best, but two assumptions seemed appropriate: (1) steady-state Isp is well known and should not be changed, and (2) Isp is likely less well known for lower duty cycles and thus smaller values of Isp. With these two guidelines, the Cassini propulsion team elected to use a single "additive constant" to all pulse-mode Isp values to force the consumption model to agree with the tank model. A multiplicative constant was not appropriate, since (1) it presumed equivalent error percentages across pulse-mode Isp values and (2) it rendered spacecraft turn and Titan flyby pulse-mode Isp to be higher than steady-state Isp! When this effort was undertaken in 2006, a value of 27.0 s for the additive constant brought the consumption and tank model curves into perfect agreement. Put another way, the best agreement between tank and consumption models through 2006 came from retroactively modeling all limit duty cycling Isp at 147 s, RWA biases at 167 s, spacecraft turns and Titan flybys at 207 s, and steady-state RCS TCMs/OTMs at 217 s. This technique was admittedly empirical, but gratifyingly it seems to have done the trick. This may be verified below in Figure 28, the final reckoning of remaining hydrazine mass vs. mission time over the twenty-year Cassini mission. Even though Isp values were last tweaked in 2006 (around mission day 3200), the agreement between consumption and tank models remained excellent for the next eleven years, over an impressive 70 kg of hydrazine usage. By the time of Cassini's plunge into Saturn, this Isp-modified throughput model suggested 132 – 34 = 98 kg of N_2H_4 was expended, or 98/132 = 74% of the total launch load including unusable hydrazine. As before, the more believable tank model indicated only 132 – 36 = 96 kg of N_2H_4 was used, or 96/132 = 73% of the launch load. Equivalently, it appears the Isp-modified consumption model overpredicted hydrazine usage by a mere 2% (98/96 – 1) vs. the tank model. Considering how little hydrazine had been used by 2006, the agreement in Figure 28 is excellent.

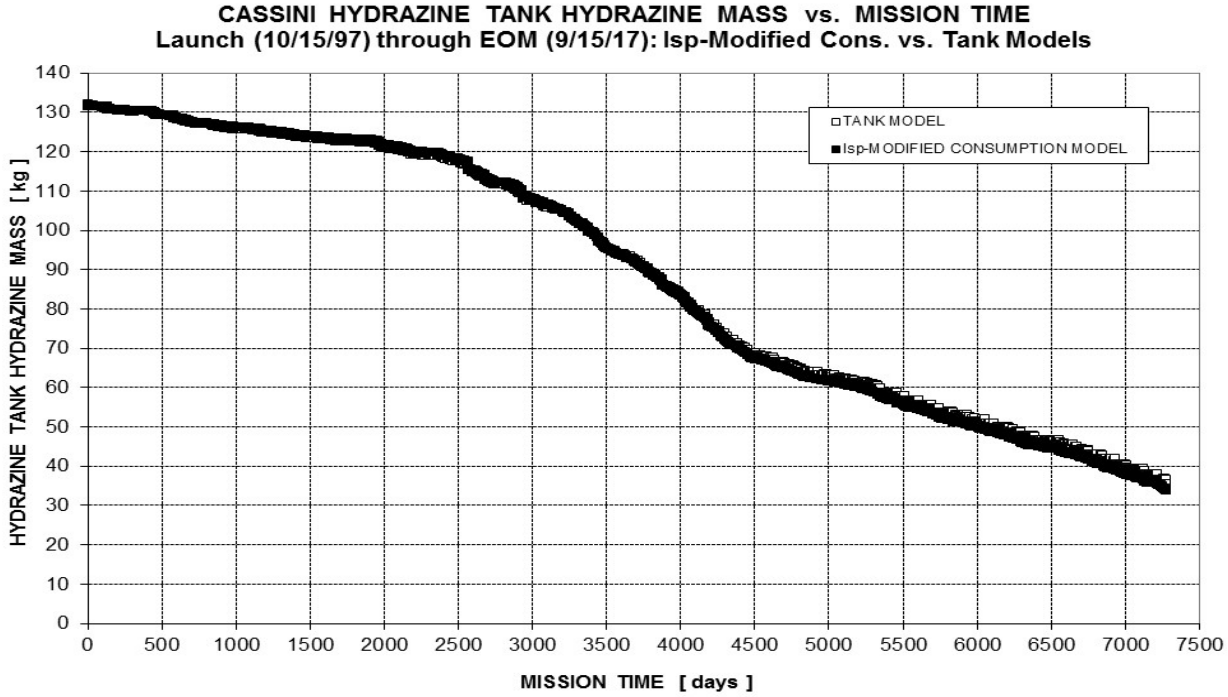

Figure 28: Cassini MTA Hydrazine Mass vs. Mission Time (Isp-Modified Consumption vs. Tank Models)

Turning now to helium budgets within the monopropellant system, the pyro-isolated helium recharge tank offered a textbook example of applying the classic helium budget method. The helium content of the recharge tank (as determined from telemetered tank pressure and temperature) is presented in Figure 29, covering the period between launch and the moments just prior to firing PV40 to "recharge" the hydrazine tank on April 10, 2006. As expected for this pyro-isolated tank, the helium content is essentially constant vs. time. This is as presumed for a pyro-isolated tank, but pleasingly this also offers additional insight into recharge tank pressure transducer drift, or

rather the lack of drift. If the recharge tank pressure transducer were drifting at the same rate as prior JPL pressure sensors over the 8.5 years spanned in Figure 29, the inferred helium mass would definitely show linear but decidedly non-zero changes, above the noise level of the data. The good news continues for Cassini's apparently drift-free pressure transducers.

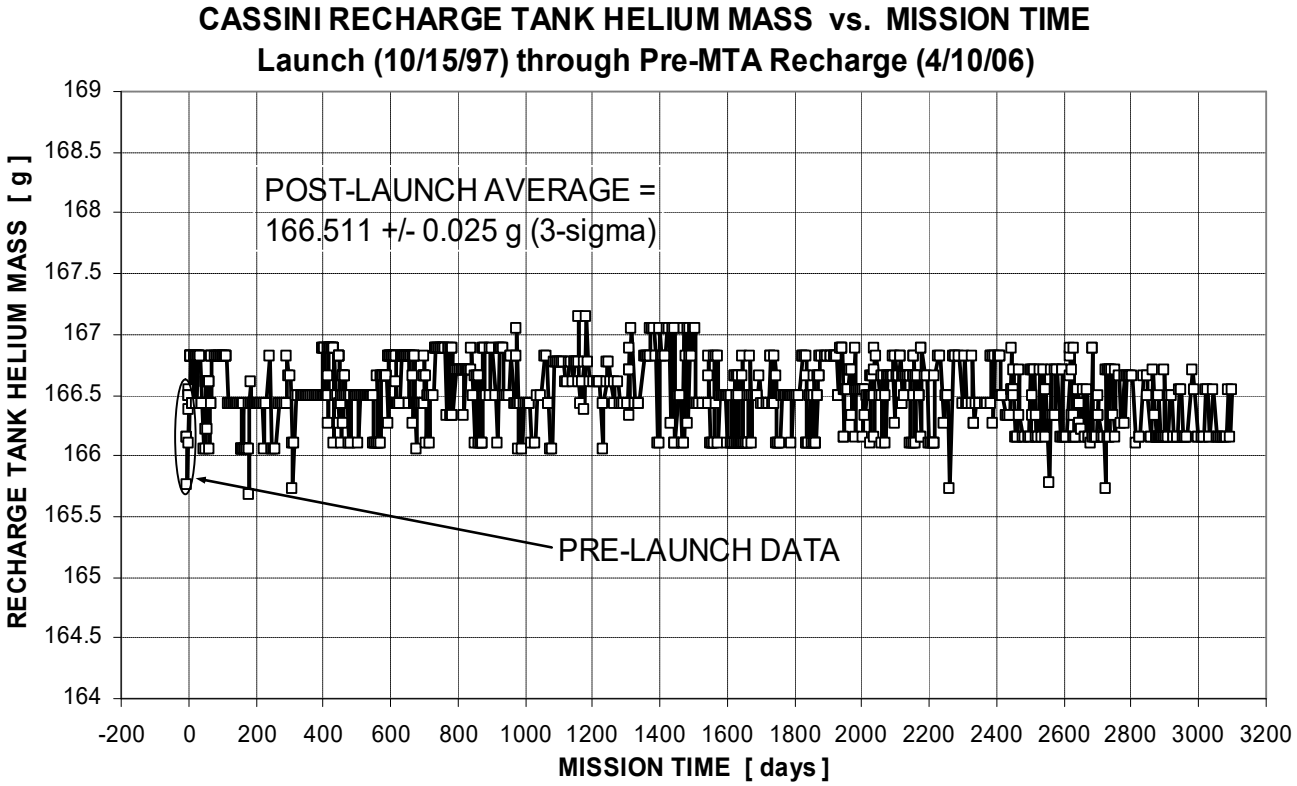

Figure 29: Cassini Recharge Tank Helium Mass vs. Mission Time (Launch through Pre-Recharge)

Figure 30 below is a bar chart showing the calculated post-launch average loaded RTA helium mass (solid horizontal line) as compared to the propellant loading analysis memo, target load, and two pre-launch calculated RTA helium mass values, one right after propellant loading and the other just before launch. Unmodeled temperature uncertainties following propellant loading render the 3σ values for the second point in Figure 30 indeterminable. The recharge tank helium load of 166.511 ± 0.025 grams (3σ) based on all data points in Figure 29 is very close to the target load of 167.78 ± 0.34 grams (3σ), but formal statistics suggest the target load was not quite met. However, the discrepancy was not large, and the hydrazine tank was slightly overloaded with helium vs. requirements (see below), so this had no consequences for mission operations. Sufficient helium was loaded into the monopropellant system to execute the prime and extended Cassini missions.

Given the vagaries of hydrazine consumption over the mission, it is not particularly useful to prepare a chart of MTA helium mass as a function of time over the mission. In fact, to generate Figures 25 and 27-28 above, the initial helium mass in the MTA was a necessary input, along with helium solubility modeling as hydrazine was expelled (admittedly, this had very small effect, given quite low helium solubility rates in hydrazine). However, it is still possible to assess helium loading masses and uncertainties in the MTA, as was done above for the RTA. Figure 31 is the MTA analogue to Figure 30, and it shows sufficient helium was loaded in the hydrazine tank to execute the Cassini mission. This is particularly true when comparing the post-launch average value of 232.5 grams vs. the propellant loading analysis value of 230 grams.

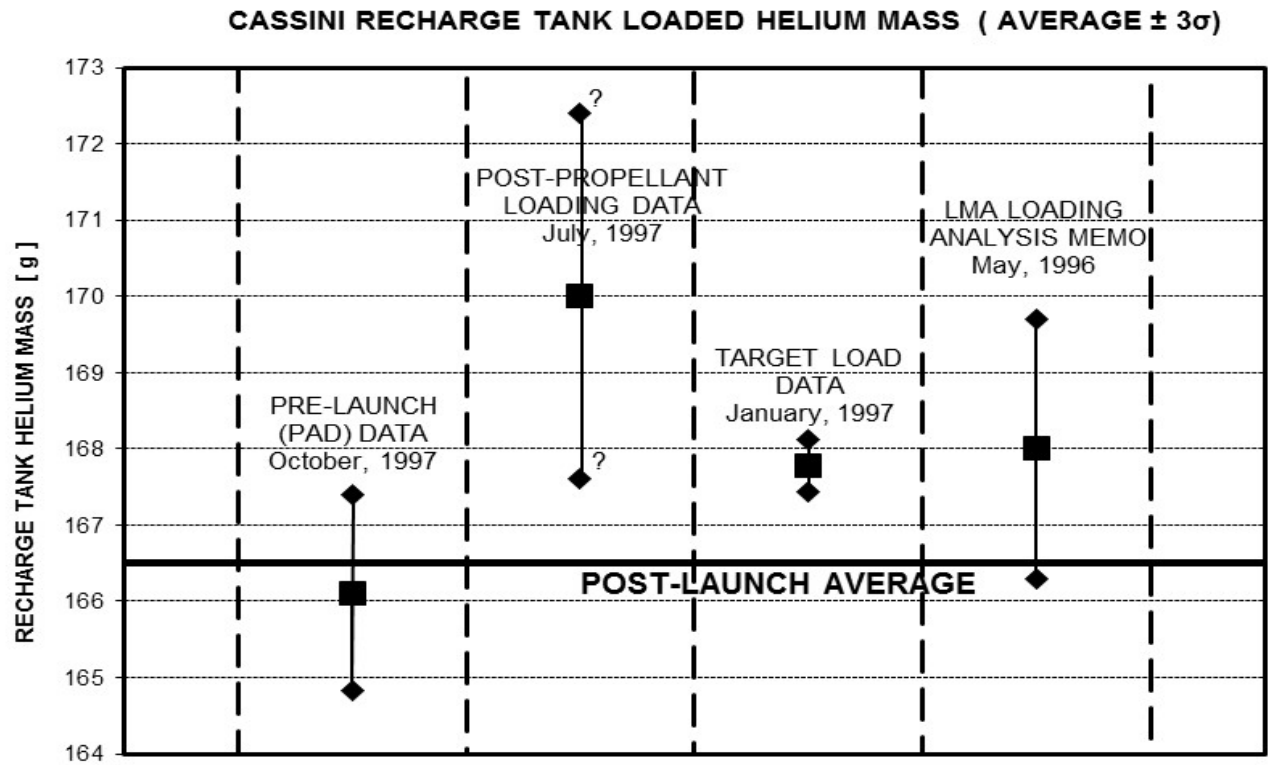

Figure 30: Cassini Recharge Tank Loaded Helium Mass Assessment

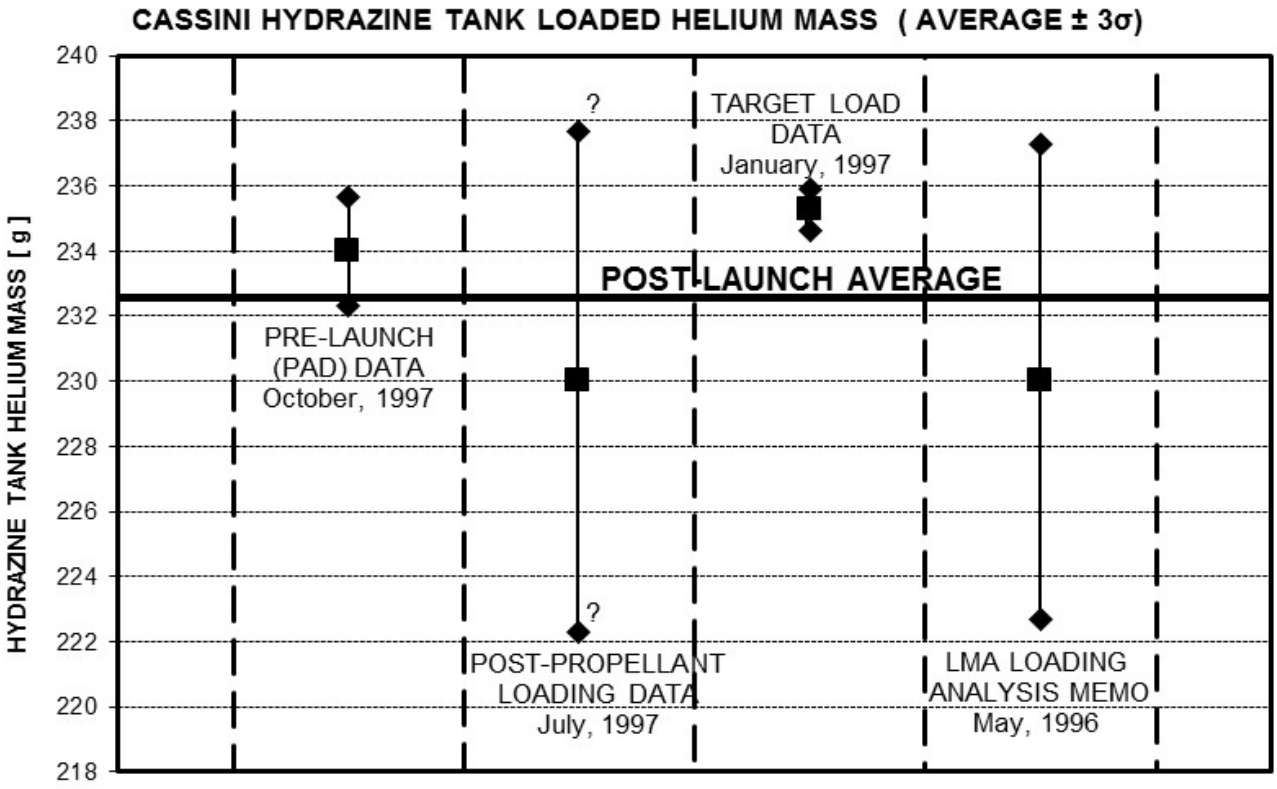

Figure 31: Cassini Hydrazine Tank Loaded Helium Mass Assessment

To broadly summarize this section, Cassini hydrazine consumption was initially overpredicted by about 19%. Accounting for flying Moog valve (MR-103H) thrusters vs. the provided Wright valve ground test model reduced errors to 16%, and a one-time empirical Isp correction to pulse-mode models many years ago decreased overall consumption modeling errors to only 2% vs. the likely more accurate tank model. Cassini's RTA pressure transducer joined the list of apparently drift-free sensors already mentioned in Section VI. Finally, cross-checks showed both the RTA and MTA were sufficiently loaded with helium to execute the Cassini prime and extended missions. It seems such a shame to have entered Saturn's atmosphere with roughly a quarter of a tank of hydrazine remaining, but bipropellant margins were far more precarious, as may be seen in the next section.

IX. Bipropellant NTO/MMH & Helium Budgets

In theory, calculating bipropellant usage is quite a bit simpler than hydrazine usage, largely because all bipropellant was consumed during steady-state main-engine TCMs and OTMs. Therefore, the considerable uncertainties associated with pulse-mode operation mentioned above were not applicable, and the best reckoning of NTO and MMH mass expenditure was obtained from bipropellant system prediction tool (bspt) reconstruction of each main-engine burn. One complication, however, was the proper accounting of helium mass within bipropellant plumbing. Specifically, some helium was expelled overboard with each main-engine burn, with helium being much more soluble in NTO than it is in MMH or N_2H_4. Occasionally, some helium was also transferred from the HTA to the BTA during regulated burns or fuel-side repressurizations. Despite these complications, bipropellant consumption errors from maneuver reconstruction were likely within a few percent. Ominously, though, final NTO and MMH masses at the time of the Saturn plunge were also in this range. It is not an exaggeration to say there were literally a few sleepless nights during the last few years of the mission, with flight team members wondering if an upcoming Cassini main-engine OTM would lead to bipropellant depletion.

Figure 32 displays remaining NTO and MMH mass during the two-decade Cassini mission, based on bspt reconstruction of all 183 main-engine TCMs and OTMs. It is very unusual to show remaining propellant mass vs. mission time on a semi-log scale, but using such a scale is the only way to convey how nerve-wracking Cassini bipropellant margins truly were. In particular, note years of operation below the 3σ unusable propellant thresholds, for both NTO and MMH! Main-engine OTMs during this time had up to a 5-10% chance of depleting either MMH or NTO, so contingency procedures were put in place to complete required mission ΔVs imminently using B-branch RCS thrusters.[13] This time in the mission was definitely not for the faint of heart, but pleasingly there were no hints of impending MMH or NTO depletion through OTM-467, the final bipropellant consumption event on December 4, 2016. Many Cassini spacecraft teams (Propulsion, AACS, Thermal, and Navigation) scrutinized late-mission OTM performance carefully, trying to discern subtle early indications of bipropellant depletion, including unexpected changes in bipropellant system temperatures, line pressures, Pc, chamber temperature, engine gimbaling, or delivered ΔV, to name a few. Thankfully, no signs of the onset of bipropellant depletion were ever observed in flight. Regrettably, though, a "burn to depletion" in the last few days of the mission was disapproved, given the high scientific value of unique Saturnian plunge data. A burn to depletion would have allowed excellent ground truth for bipropellant consumption models and errors, but it would have incurred some risk to Cassini's precious final bits of *in situ* science.

Turning now to bipropellant system helium budgets, it may be possible to further assess NTO and MMH mass remaining by attempting to bean-count every atom of helium transferred from the HTA to the BTA or jettisoned to space during 183 REA-A firings. Helium budget cross-checking calculations were developed for the Galileo mission to Jupiter (see Ref. 3) and applied to Cassini early in the mission (see Ref. 6). Figure 33 is the "classic" helium budget plot, showing total bipropellant system helium mass as a function of mission time, including the individual tank helium mass curves which add up to the total mass. This single plot summarizes Cassini's bipropellant helium mass history by tank, along with the HTA-to-BTA helium transfer history. Starting with the bottom curve and working up, inferred helium masses in the HTA, NTO tank ullage, MMH tank ullage, and NTO/MMH solution are "stacked up" in Figure 33. Therefore, the top curve should start at a value equal to the launch load of bipropellant system pressurant, while the second curve from the top, the sum of HTA helium and BTA ullage helium masses, should remain nearly constant over the mission. The data in Figure 33 are all as expected, but by examining the curves in Figure 33 at much higher resolution, further details may be gleaned for (1) Cassini's nine HTA-to-BTA helium transfer events, (2) helium solubility rates and time constants into NTO and MMH, and (3) potential bipropellant consumption errors.

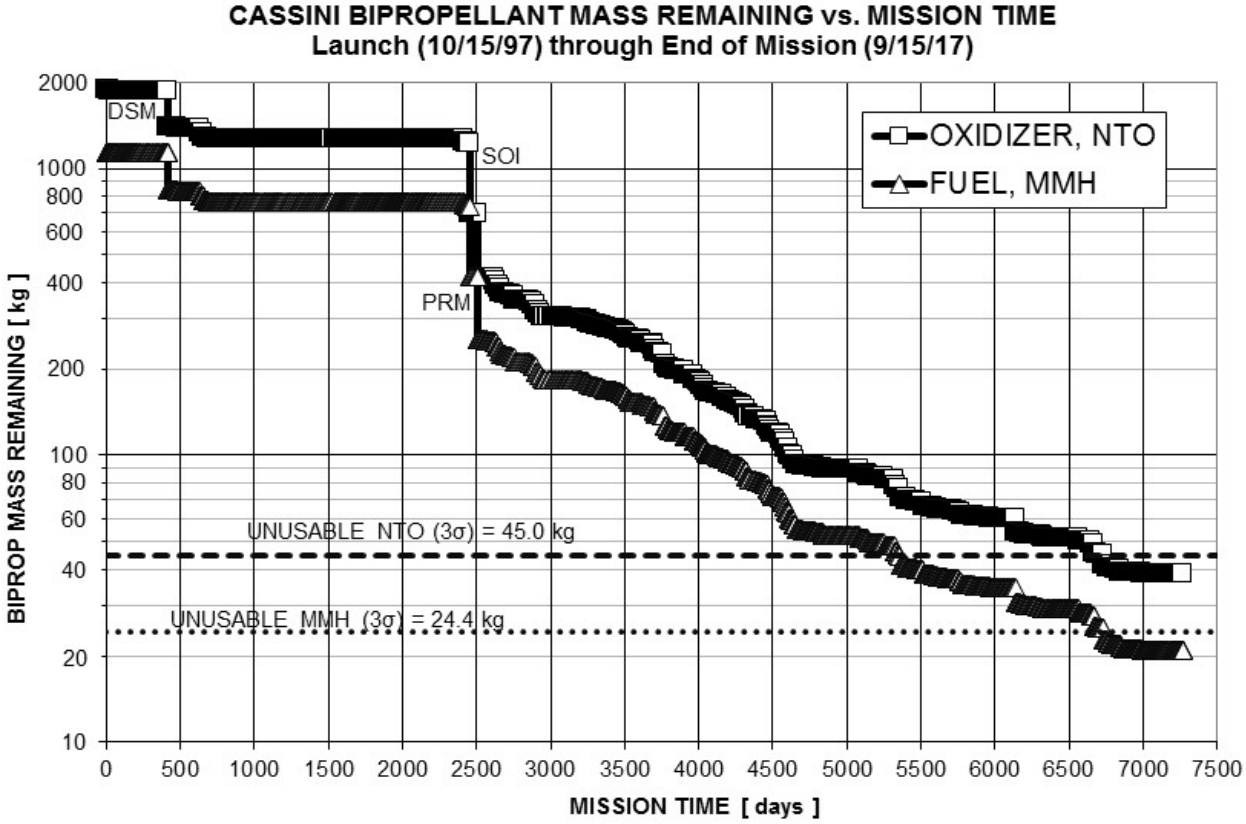

Figure 32: Cassini NTO/MMH Mass Remaining vs. Mission Time (Including Unusable, Semi-Log Scale)

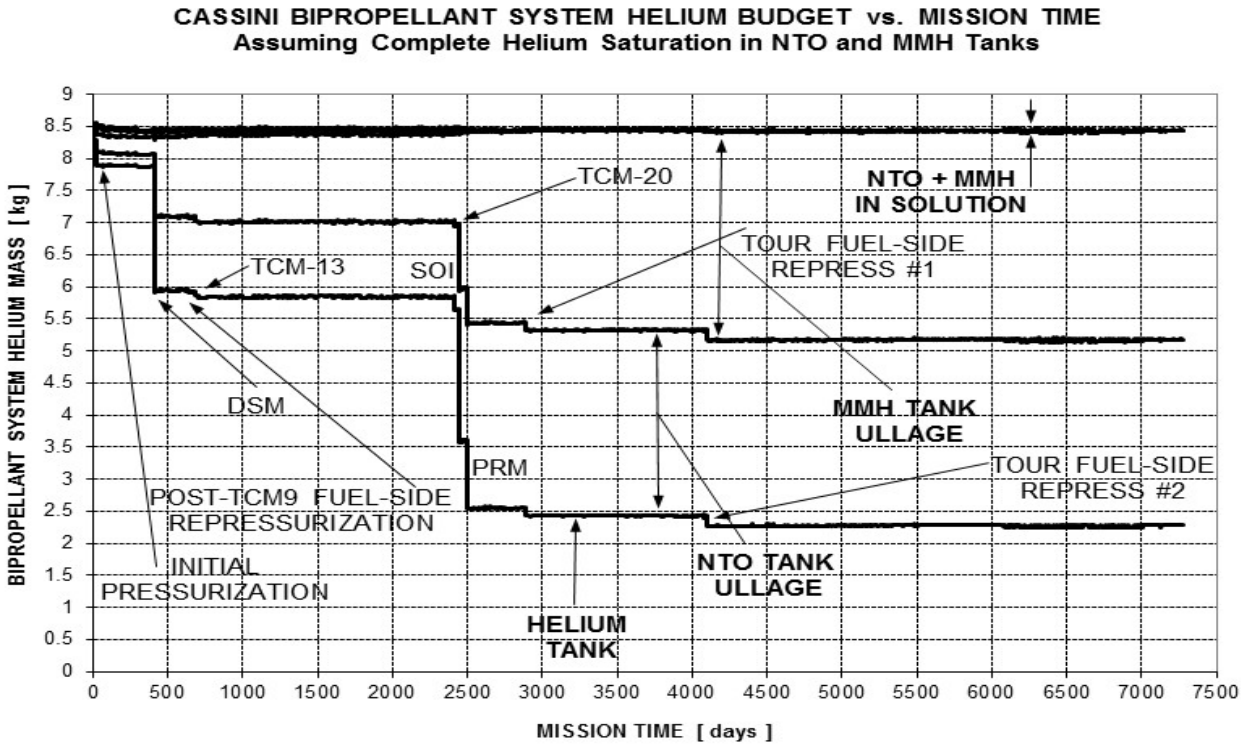

Figure 33: Cassini Bipropellant System Helium Budget (HTA/BTA Helium Masses) vs. Mission Time

Focusing on just the bottom curve in Figure 33, a "decluttered" plot of the HTA inferred helium mass over the mission is presented below as Figure 34. It could not have been simpler to calculate the mass of helium in the HTA; essentially, the Beattie-Bridgeman equation of state was solved with tank volume and telemetered tank pressure and temperatures as inputs. The data of Figure 34 resemble step functions, as expected for the case of a finite number of short-duration helium transfer events during the Cassini mission. Less obvious, the periods of time in Figure 34 between pressurization events may be investigated for possible HTA pressure transducer drift, with some windows many years in duration. Not surprisingly, HTA helium budgets during these periods of stasis had basically zero slope. It seems yet another Cassini pressure transducer exhibited drift-free performance over nearly two decades.

Theoretically, with PR1 open for nearly the entire mission, gross LV10 leakage might have been observable in Figure 34, but realistically the coarse DN resolution of 16.8 psia/DN in the HTA pressure transducer would incur large error bars for any such analysis. In contrast, changes in PR1 outlet pressure (see Figure 8) during the quiescent periods from Figure 34 ended up offering exquisite insight into LV10 helium diffusion, thanks to finer PR1 outlet pressure DN resolution (1.88 psia/DN) and especially the tiny volume between PR1 and LV20/30 or CV20-23/CV30-33. Helium diffusion and potential leakage through LV10 will be explored in detail in the next section.

There were nine HTA-to-BTA helium transfer events during the Cassini mission, namely (1) initial pressurization, (2) DSM, (3) post-TCM9 fuel-side repressurization, (4) TCM-13 half-regulated maneuver, (5) TCM-20, (6) SOI, (7) PRM, (8) fuel-side repressurization #1 of tour, and (9) fuel-side repressurization #2 of tour. The step functions caused by these events are all readily visible in Figure 34, though there clearly were large differences in transferred helium mass depending on the event. The only surprise in Figure 34 is an apparent 0.1 kg decrease in helium mass immediately after launch. This initially caused some concerns about helium leakage, but a more mundane explanation was eventually determined. Pre-launch, the HTA was in thermal equilibrium, with no thermal gradients across the tank. In contrast, right after launch the top of the HTA cooled ~ 4°C (due to exposure to space) while the temperature at the bottom of the HTA increased ~ 4°C (due to RTG waste heat). On average, there was no change in HTA bulk temperature, so HTA pressure remained unchanged. However, ground accounting of HTA helium mass did change, because both temperature transducers on the HTA were mounted to the tank bottom. One lesson learned from this minor post-launch surprise was to distribute temperature sensors more evenly around tanks with thermal gradients.

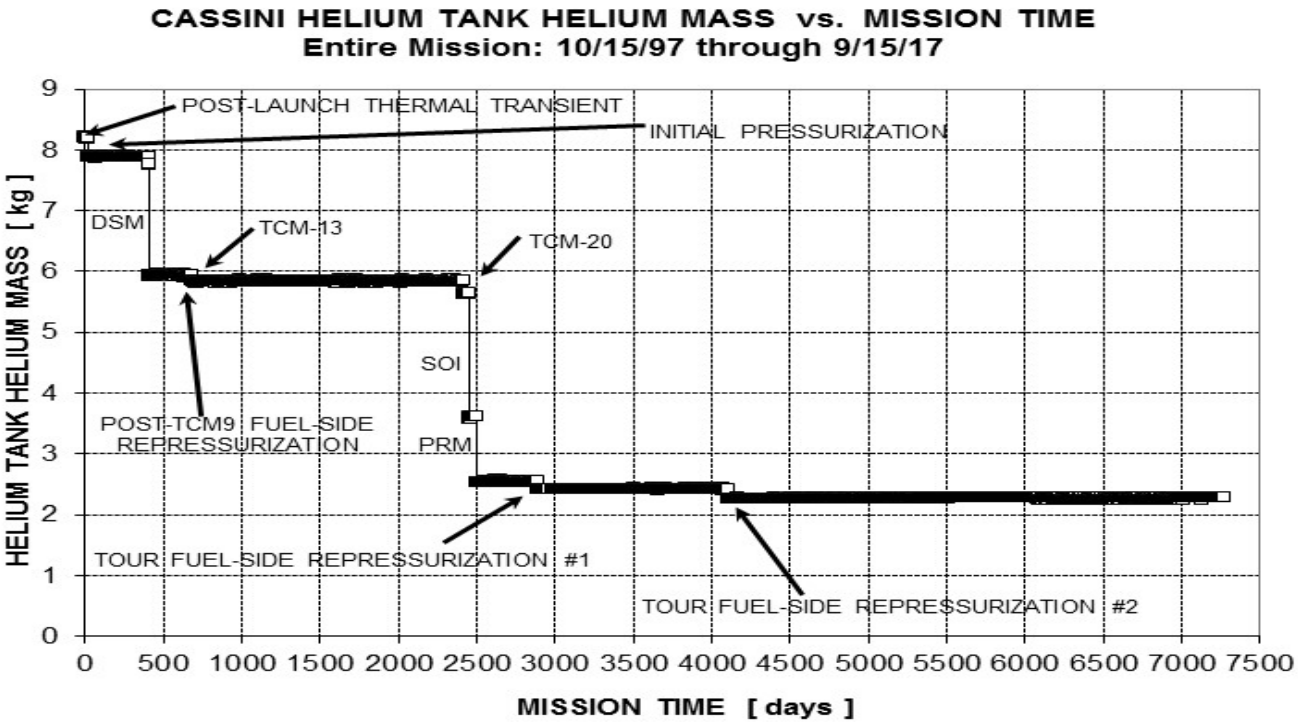

Figure 34: Cassini Bipropellant System Helium Tank Inferred Helium Mass vs. Mission Time

The difference between the bottom two curves in Figure 33 represents the inferred helium mass in the NTO tank ullage volume as a function of mission time. It is difficult to mentally interpret the difference between two curves by examining Figure 33, however, so this quantity is presented explicitly below in Figure 35. As with Figure 34, this cleaned-up version immediately shows the expected behavior for NTO tank ullage helium mass. The only times helium was transferred from the HTA to the NTO tank were during initial pressurization and Cassini's four fully regulated main-engine maneuvers (DSM, TCM-20, SOI, and PRM). Therefore, as expected, Figure 35 essentially shows five step function "jogs" over the twenty-year Cassini mission.

One deviation from the "stair-step" curve expected in Figure 35 is the interesting behavior resembling a decaying exponential after initial pressurization. Helium budget modeling assumes complete helium saturation in NTO and MMH solution (i.e., it is an equilibrium calculation), so it is not surprising to see the extended effects of helium slowly going into NTO solution after initial pressurization, particularly given the large NTO tank pressure change of 114 psia to 257 psia at that time. From Ref. 6, the first-order time constant for helium saturation in NTO turned out to be 82 days, similar to Mars Observer's NTO/MMH saturation time constant (on the ground) of about 50 days. Helium solubility levels in NTO agreed with ground models within about 20% as well. Further details are available in Ref. 6.

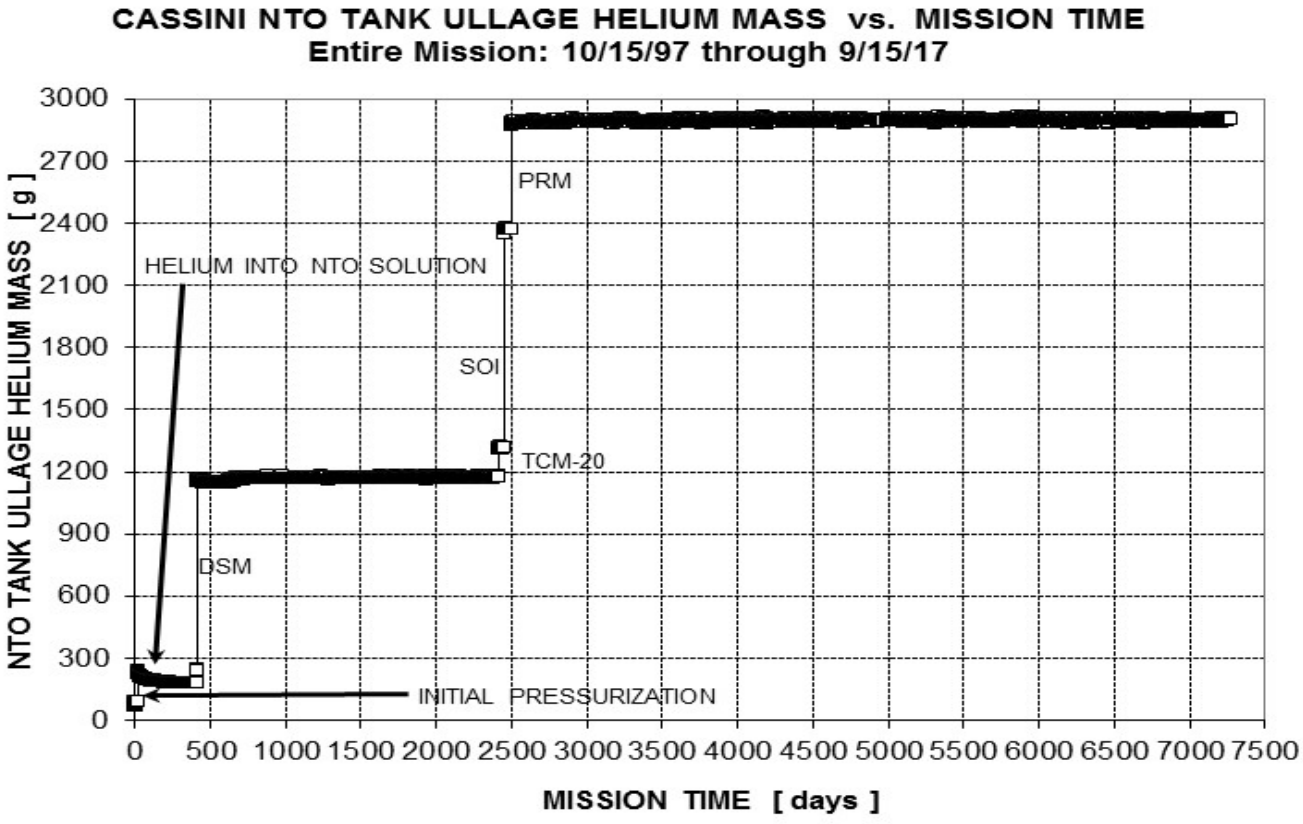

Figure 35: Cassini Bipropellant System NTO Tank Inferred Ullage Helium Mass vs. Mission Time

During the periods between NTO-side helium transfer events, in theory it might be possible to calculate NTO mass remaining, given nearly constant NTO total (ullage plus solution) helium mass values between pressurizations. This is particularly true given pyro-isolation of the NTO tank for most of these periods. Put another way, NTO consumption errors might be discernible in Figure 35 if inferred NTO helium mass values did not appear to stay constant. Detailed examination and scale magnification of the periods between pressurization events in Figure 35 showed near-zero slopes within high noise level data ranges. This suggests NTO mass modeling was likely quite

good, but large error bars prevented further quantification of actual NTO consumption during these blowdown periods, unfortunately.

The MMH analogue to Figure 35 is presented below as Figure 36. In addition to the five helium transfer events already mentioned for the NTO tank, helium was also transferred from the HTA to the MMH tank during three fuel-side-only repressurizations (post-TCM9 and two during tour) and throughout TCM-13, Cassini's only half-regulated maneuver. As expected, Figure 36 essentially shows nine step function "jogs" over the twenty-year Cassini mission for these MMH pressurization events. Similar to before, the extended effects of helium slowly going into MMH solution after initial pressurization are apparent. MMH tank pressure increased from 103 psia to 256 psia at initial pressurization, though helium is much less soluble in MMH than it is in NTO (hence the smaller "signal" in Figure 36 vs. Figure 35). From Ref. 6, the first-order time constant for helium saturation in MMH was around 81 days, almost identical to the 82-day value for the NTO tank. Helium solubility levels in MMH agreed with ground models to less than 10%. Additional details regarding early mission helium solubility modeling are available in Ref. 6.

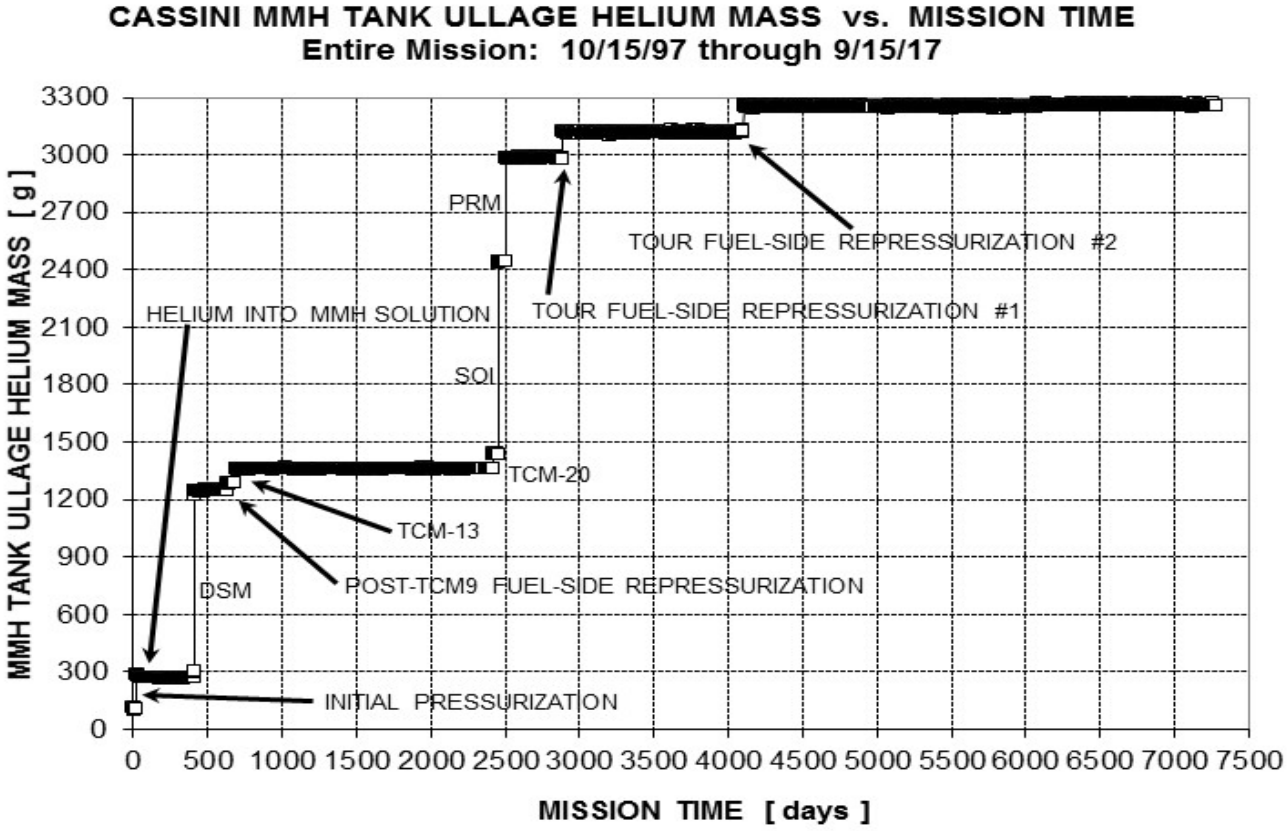

Figure 36: Cassini Bipropellant System MMH Tank Inferred Ullage Helium Mass vs. Mission Time

As before, in theory it might be possible to determine MMH usage errors by examining the flat periods in Figure 36 to look for variability. This is complicated a bit by the potential for helium flow through (admittedly closed but forward-relieving or possibly leaking) LV30 and CV30-33 between pressurization events (more on this in the next section). The good news from Figure 36 is MMH usage during blowdown maneuvers was not grossly mismodeled, given near-zero slopes for periods between pressurization events in Figure 36. However, as it was for the NTO tank, large error bars made it impossible to assess MMH consumption errors during blowdown maneuvers. In addition, NTO and MMH consumption errors during the Cassini mission were dominated by uncertainties in DSM, SOI, and PRM bipropellant usage, periods with non-constant helium mass in NTO and MMH tank ullages.

From Figures 34-36, large amounts of helium were transferred from the HTA to the BTA during DSM, SOI, and PRM. Lesser amounts were transported downstream during other pressurization events; however, as expected, the total helium mass remained essentially constant across all pressurization events, as may be verified approximately in Figure 33. This can be confirmed more rigorously by tabulating the helium masses transferred during Cassini's nine bipropellant pressurization events (see Table 15), using the fine detail available from Figures 34-36, curve fitting average values of helium mass just before and after pressurization events, etc. In all nine cases, the amount of helium transferred as calculated by the decrease in helium tank (HTA) helium mass agreed very well with the amount calculated by the increase in bipropellant tank (BTA) helium mass. These results were as expected (within uncertainties) for the case of zero external helium leakage. Also as expected, errors were generally higher for smaller helium transfer events, such as the two fuel-side repressurizations of tour. In summary, Table 15 shows excellent agreement for the conservation of helium mass across all Cassini bipropellant pressurization events.

Table 15: Cassini Bipropellant Pressurization Event Helium Mass Transfer Summary

Pressurization Event	Date [M/D/Y]	NTO ΔHe MASS [kg]	MMH ΔHe MASS [kg]	BTA ΔHe MASS [kg]	HTA ΔHe MASS [kg]	BTA/HTA ERROR [%]
Initial Press.	11/9/97	0.139	0.181	0.320	-0.309	-3.4
DSM	12/3/98	0.916	0.948	1.864	-1.831	-1.8
Post-TCM9	7/13/99	0	0.352	0.352	-0.385	+9.4
TCM-13	9/1/99	0	0.731	0.731	-0.769	+5.2
TCM-20	5/27/04	0.145	0.078	0.223	-0.204	-8.5
SOI	7/1/04	1.055	1.004	2.054	-2.036	-0.9
PRM	8/23/04	0.514	0.540	1.054	-1.055	+0.1
Fu Repress #1	9/12/05	0	0.133	0.133	-0.118	-11.3
Fu Repress #2	1/8/09	0	0.137	0.137	-0.158	+15.3

One final attempt to assess bipropellant consumption errors was undertaken using helium budget data. As was mentioned previously, the only change in total bipropellant helium mass over the mission should be the loss of helium in NTO and MMH solution expelled during main-engine TCMs and OTMs. Otherwise, total helium mass should remain constant, simply shifting from the HTA to the BTA a handful of times over the mission. Great pains were taken to model helium quantities in NTO and MMH solution during the mission, including assessing helium saturation quantities and time histories after each pressurization, scaling the initial pressurization time constant values of 81-82 days by a representative length scale for diffusion (the distance from the top of a BTA tank to the liquid/gas interface in the tank), accounting for differing solubility levels as a function of propellant liquid temperature, and decrementing the appropriate amount of helium mass for each main-engine burn. The results of this painstaking effort are presented below in Figure 37, a time history of the total jettisoned helium mass in NTO and MMH solution caused by 183 REA-A firings. Gratifyingly, the total helium mass expelled to space from Figure 37 was 129.5 g, within 2% of the total helium mass modeled to be in NTO and MMH solution after initial pressurization, 127 g. Though both calculations use the same helium solubility models, large variations in liquid propellant temperature over the mission could have caused significant discrepancies between these two quantities.

Armed with decent knowledge of the helium lost to space over the mission, this quantity may be subtracted from the well known bipropellant system total helium launch load of 8.550 kg. This difference is the total bipropellant helium mass remaining, as a function of mission time, and thus it should be identical to the top curve in Figure 33. On a vastly expanded vertical scale, these two independent values for total system helium mass are presented below in Figure 38. Unlike previous helium mass comparisons between pressurization events in Figures 35-36, the difference between the curves is now statistically significant, well above the noise level of the helium budget data. Even with uncertainties in helium solubility in NTO and MMH, by the end of mission the difference between the curves in Figure 38 was greater than 100 grams, approaching solubility levels themselves. Since 100% errors are very unlikely in helium solubility modeling (especially given the excellent agreement in helium solubility modeling during initial pressurization mentioned above), the most likely explanation for the discrepancy in Figure 38 is mismodeling NTO and/or MMH consumption. Fortunately, it is in the direction of having more NTO and/or MMH than modeled based on bspt reconstructions. Regrettably, though, it is impossible to determine from Figure 38 which bipropellant constituent might have been more plentiful than modeled. This will be discussed further below.

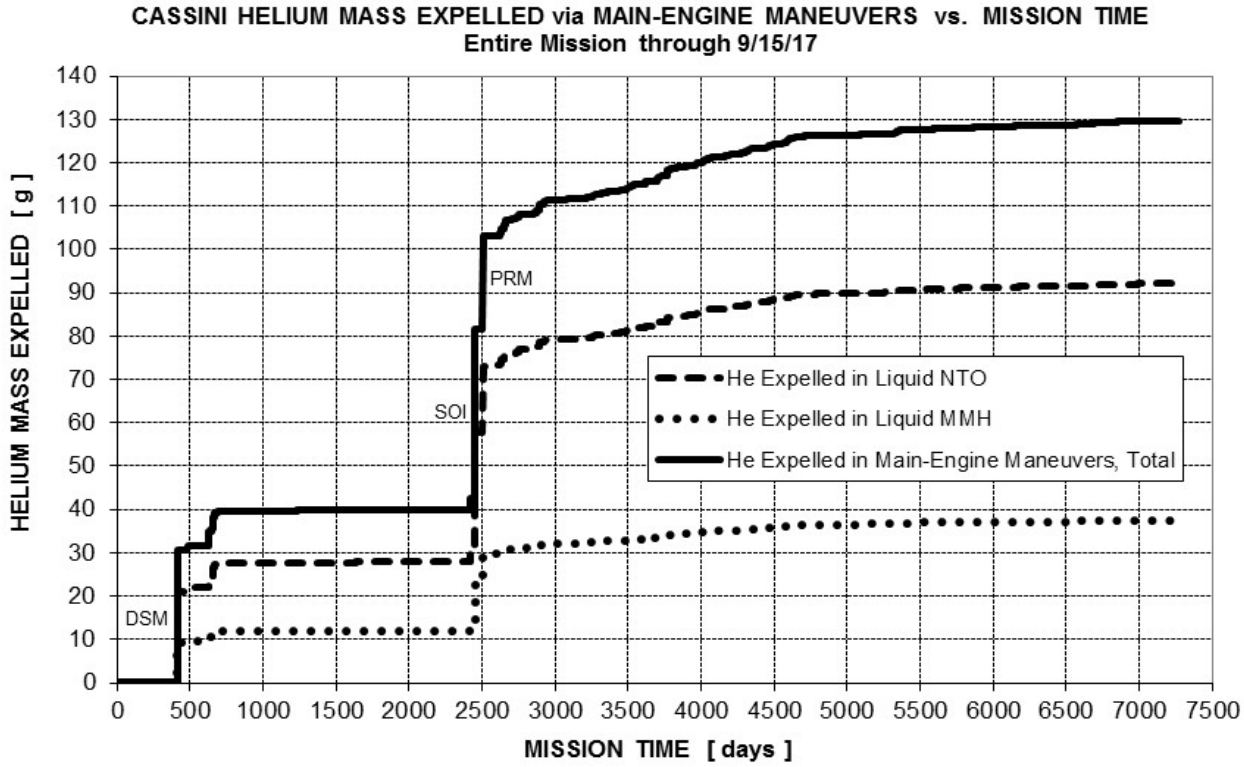

Figure 37: Cassini BTA Helium Mass Expelled During REA-A TCMs/OTMs vs. Mission Time

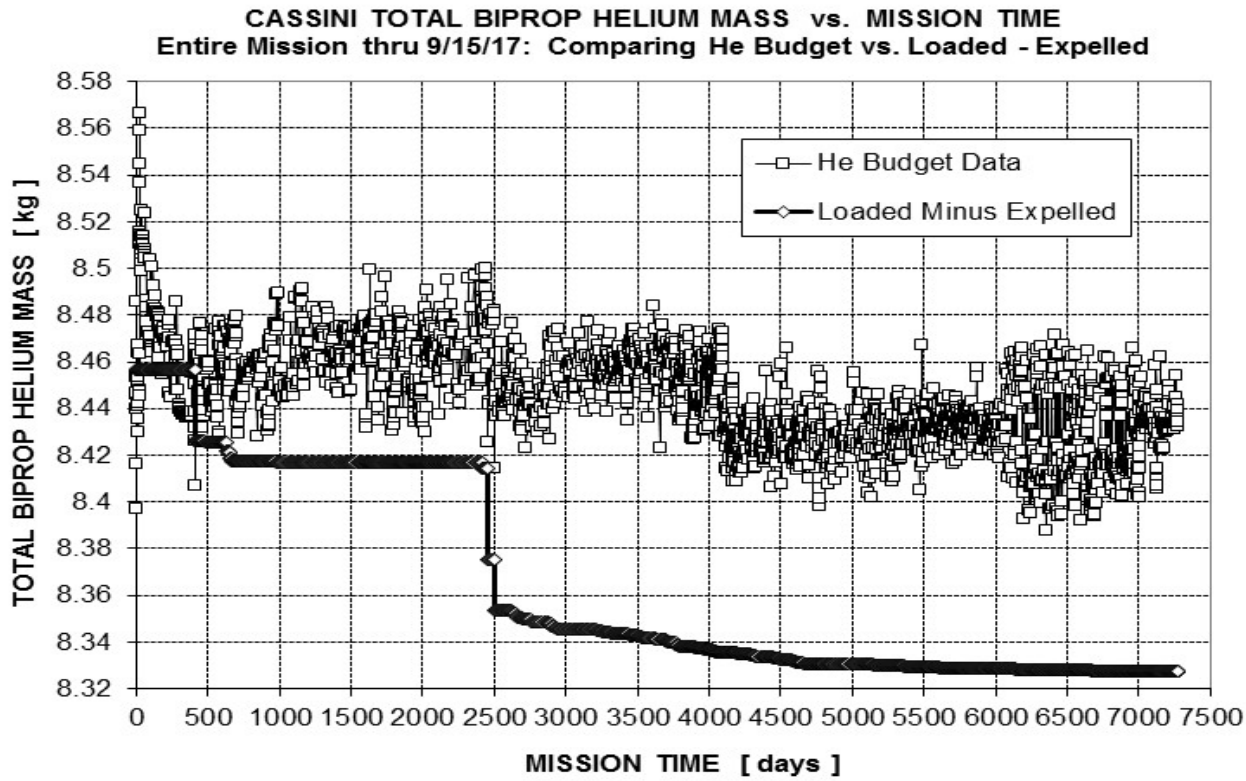

Figure 38: Total Biprop Helium Mass vs. Mission Time (He Budget & Loaded Minus Expelled)

If helium solubility and expulsion were modeled perfectly, along with NTO/MMH consumption, the top curve in Figure 38 would match the bottom curve exactly. One way to explore the disagreement between the two data sets is to plot their percentage difference vs. the percentage of bipropellant used, rather than vs. mission time (see Figure 39). Bipropellant consumption was not linear vs. mission time, but if the discrepancy in Figure 38 were due to usage errors, in theory this error may well have built up linearly with propellant usage. Even though the noise level is high, the linear fit correlation of $R^2 = 0.73$ in Figure 39 is not bad. The data of Figure 39 are "unusually" bunched together at different positions along the x-axis, but this is as expected, given large bipropellant expenditures in DSM, SOI, and PRM (readily visible below as the three "chasms" along the x-axis, moving from left to right).

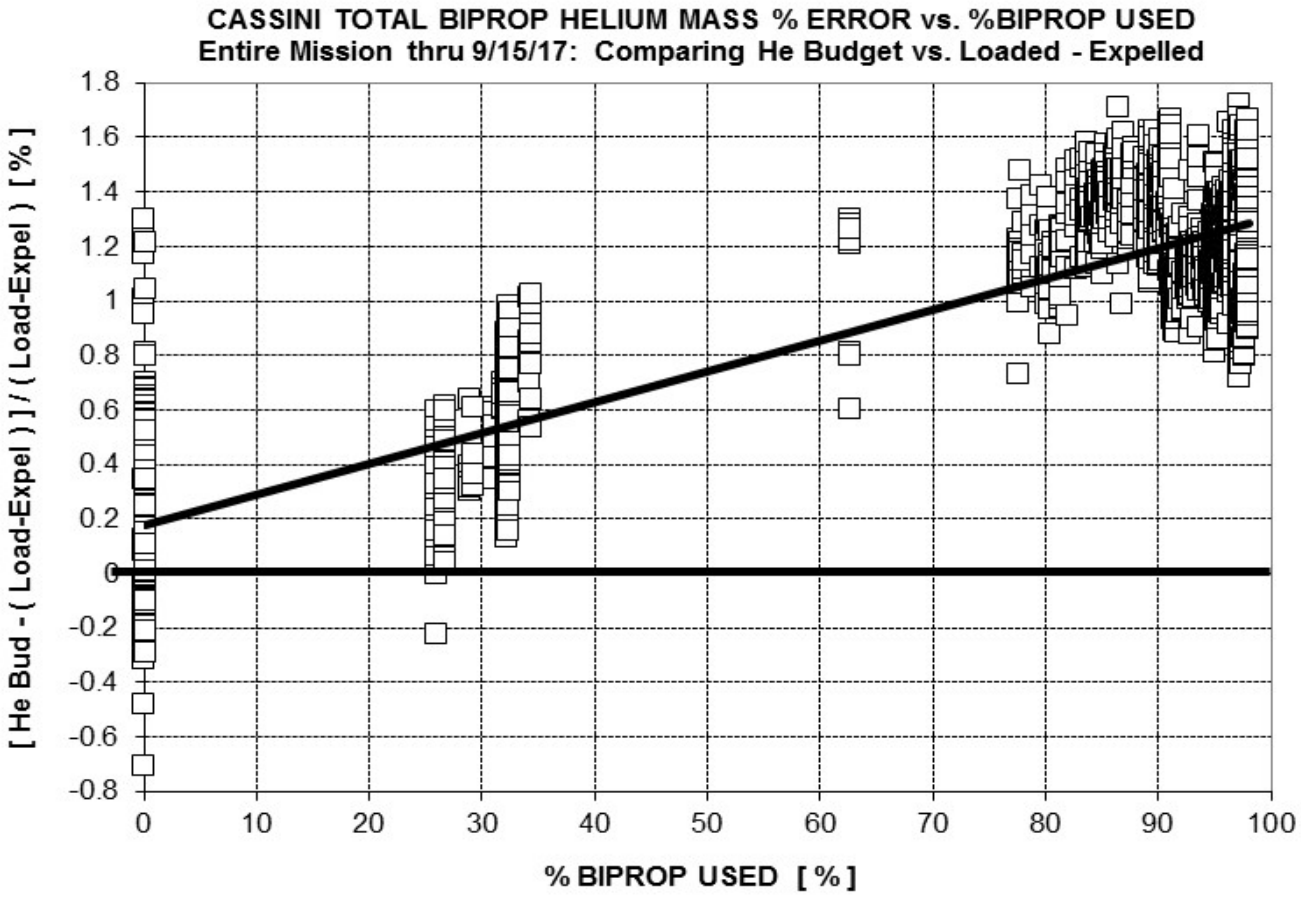

Figure 39: Cassini Total Bipropellant Helium Mass %Error (He Budget vs. Loaded – Expelled) vs. %Biprop Used

As mentioned previously, even though Figures 38-39 suggest less NTO and/or MMH were used than modeled in bspt reconstructions, it is impossible to determine which propellant constituent may have been overcounted. In a Gedanken experiment, though, one could calculate the end-of-propellant (EOP) "extra" mass in either the NTO or MMH tank which would "flatten out" the curve fit in Figure 39. These values turned out to be rather substantial, about 65.3 kg of extra NTO or 34.5 kg of extra MMH. These represent 3.49% and 3.05% of the launch loads of NTO and MMH, respectively. Of course, an infinite number of combinations of NTO and MMH usage errors also would render the curve in Figure 39 to be flat. In fact, the problem is essentially linear, so it is possible to determine the linear combination of NTO and MMH mass errors which would bring helium budget data in line with loaded minus expelled data. These ranges of bipropellant mass errors are presented below in Figure 40. Note the two first quadrant boundary "anchor" points of 34.5 kg of extra MMH with zero NTO mass error and 65.3 kg of extra NTO with zero MMH mass error. Connecting these points and linearly extrapolating to the second and fourth quadrants show the possible range of NTO/MMH mass error ordered pairs which would yield a horizontal curve fit in Figure 39.

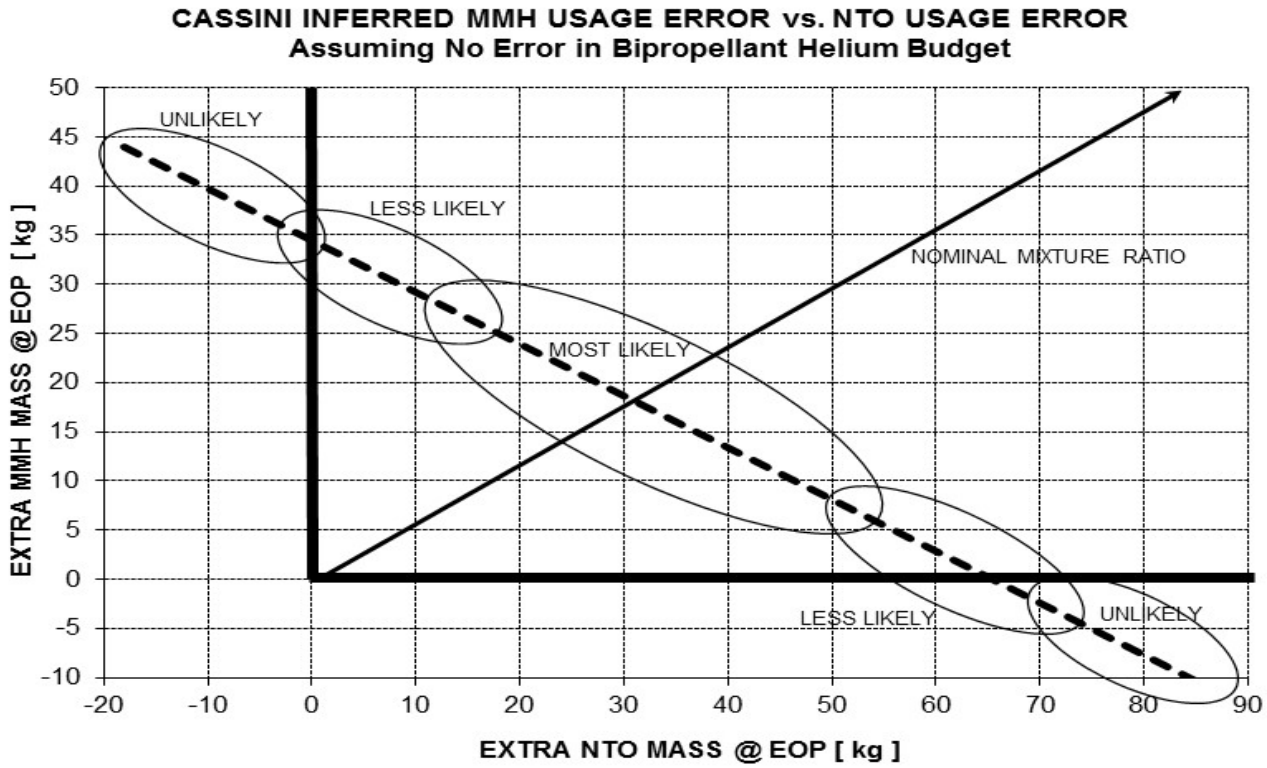

Figure 40: Inferred Cassini NTO Usage Error vs. Inferred MMH Usage Error (if Biprop Helium Budget is Correct)

Even though mixture ratio errors are certainly possible in modeling bipropellant consumption, it is unlikely they would be too excessive, and they should be distributed normally around the nominal mixture ratio of 1.6525. As such, the most likely point from Figure 40 is the intersection of the excess bipropellant line with the nominal mixture ratio line. At this point, excess modeled bipropellant usage would have only been 30.4 kg of NTO and 18.4 kg of MMH, a mere 1.63% of the launch loads of 1869 kg of NTO and 1131 kg of MMH. Errors of this magnitude are definitely possible, but uncertainty remains with respect to actual NTO and MMH mass usage. It is less likely to have ended up near the x-axis or y-axis in the first quadrant in Figure 40, and even more unlikely to be in the second or fourth quadrant. Only in the latter case would actual bipropellant margins have been worse than modeled in bspt reconstruction. Despite the lack of absolute knowledge, given the large range of ordered pairs in Figure 40, it did help assuage fears that bipropellant depletion during the mission was imminent. Knowing bipropellant margins were scarce, the flight team was open to any potential good news regarding the chances of surviving until the end of mission.

During the last few years of Cassini's flight, two more independent efforts were undertaken to assess bipropellant consumption, given its potential to undermine mission completion. Unfortunately, both of these assessments suggested there was less bipropellant (particularly MMH) remaining than modeled, at odds with helium budget results, though with large error bars. Given tight margins for finishing the mission, independent propulsion experts within JPL assessed R-4D consumption errors in bspt, based on more intimate familiarity with the actual Cassini main-engine ground tests. Fortunately, the zeroth-order conclusion was bspt predicted in-flight NTO and MMH consumption quite accurately, well within original requirements. However, the most likely values from this outsider "scrub" of as-flown bipropellant usage suggested there was perhaps 2.7 kg less NTO and 5.9 kg less MMH remaining than previously modeled. Even though these numbers were small (only 0.14% and 0.52% of the launch loads, respectively), unfortunately they were in the bad direction. Moreover, MMH margins were a bit lower than NTO margins already, so this just made matters worse. Most of the correction suggested by independent propulsion experts was due to more accurate modeling of mixture ratio shifts at increased propellant inlet temperatures, ones not modeled as accurately in bspt. The Cassini project's hopes were thus dashed that there was additional "hidden"

bipropellant margin to be found by sharpening the pencil. It was more likely than ever bipropellant depletion was a very real possibility before 2017.

The final effort to independently determine remaining NTO and MMH mass was extensive, rather bold, and a first for JPL. Given its success on other missions, including JPL-led contractor missions such as Stardust, a Propellant Gauging Test (PGT) was proposed and eventually implemented on Cassini over seventeen days in October of 2012. Essentially, PGT was an attempt to subject the NTO and MMH tanks to the largest thermal transients (heating and then cooling) safely possible, measure temperature response during the lengthy cooling period, and then infer liquid NTO and MMH masses remaining based on thermal capacitance modeling. A scientifically quiet period at solar conjunction was selected for execution, and the spacecraft performed excellently during sizable thermal transients in both directions. PGT tank temperature changes up to $\Delta T = 12°C$ or so were accomplished via heater cycling and main-engine cover actuation. Even though this ΔT is rather modest compared to levels attainable on other spacecraft, it was deemed sufficient for gleaning possible useful information, all while ensuring spacecraft health and safety.

Unfortunately, initial PGT results were a bit discouraging, to put it mildly. Even though the inferred remaining NTO mass calculated from PGT was close to expectations, the best fit for remaining MMH mass from PGT suggested there was half as much MMH available as previously thought! If PGT inferred masses were correct, the prospects for completing the Grand Finale were quite bleak. Modeling efforts continued for another 1-2 years after PGT, adding additional layers of complexity, including (1) CFD modeling of the NTO/MMH liquid surface locations, (2) DN toggle curve-fitting to temperature decay curves, and (3) error bar analyses. Fortunately, final results were a bit more palatable, but still showed 2 kg less MMH remaining than modeled via bookkeeping. For the NTO tank, final PGT results were quite disparate, but this was likely due to unmodeled NTO liquid vaporization during PGT thermal transients. Figure 41 below charts the (usable mean) total remaining NTO and MMH masses as of the PGT epoch (October, 2012), not including deterministically unusable NTO and MMH. It was at least comforting that bookkeeping and PGT results were consistent within their respective 3σ uncertainties, but the project's hopes were once again crushed for "finding" excess bipropellant vs. the bookkeeping model, at least for the more critical constituent, MMH.

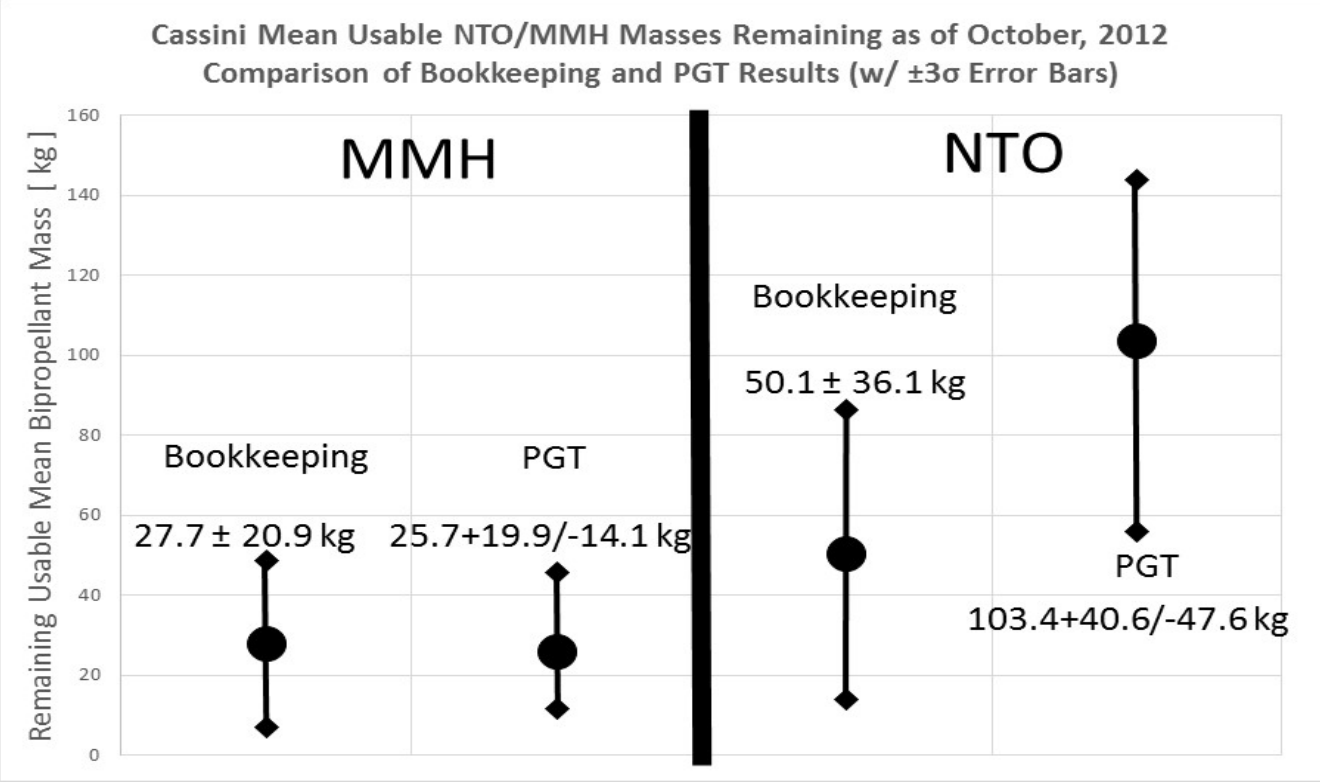

Figure 41: Cassini Mean Usable NTO/MMH Masses as of October, 2012 (Bookkeeping vs. PGT)

Cassini was never designed for PGT; for a first JPL-internal attempt at thermal capacitance modeling to determine remaining propellant, this effort should be applauded, even if the results were less than comforting. Some lessons learned from Cassini's PGT experience include (1) coarse temperature transducer resolution (0.75°C/DN) led to large error bars for the relatively modest 12°C total excursion, (2) lost Solar Thermal Vacuum (STV) transient thermocouple data prevented characterization of tank radiant surroundings, and (3) the lack of temperature sensors on the cylindrical structure surrounding the BTAs hindered proper assessment of the boundary conditions. Cassini thermal engineers provided a number of recommendations for designing future spacecraft with PGT in mind, including (1) placing temperature sensors at tank boundaries and along significant heat paths, (2) providing at least two flight transducers on the wetted portion of the propellant tank, (3) using finer DN sensors and/or building in the capability for more direct thermal transient injection (e.g., with appropriately sized tank heaters—more on this below), (4) conductively isolating tanks from the structure, (5) radiatively isolating tanks a bit with multilayer insulation, and (6) correlating appropriate thermal models during STV tests and over the mission in preparation for PGT execution(s) in flight.

Given the propensity for current missions to operate within unusable propellant margins during their mission extensions, it may be useful for future projects to provide the capability to more accurately gauge in-flight propellant usage (e.g., by using flowmeters). This system trade should be done early, but if superior accuracy is possible via other gauging techniques, it may save a lot of post-launch work and consternation for the flight team. Despite all of the challenges, it was the author's honor and privilege to serve as "human gas gauge" for the Cassini mission for a couple of decades. Running on fumes for years, with the nearest gas station 900 million miles away, Cassini's lengthy mission and valiant operation well within the margins certainly embodied JPL's new mantra to "dare mighty things."

X. Cassini Bipropellant Pressurization System Analyses

Compared to other JPL missions and particularly commercial satellites, Cassini's complex propulsion system was rather richly instrumented with pressure and temperature telemetry measurements. Given nearly twenty years of mission operations, there were occasionally opportunities to delve into propulsion component performance in great detail. There were many reasons for doing this, including improving in-flight propulsion system performance, monitoring health and safety at the component level, discovering trends before they became troublesome, and even providing rare in-flight performance data for propulsion component manufacturers. This section summarizes component performance for the bipropellant portion of the CPMS upstream of the NTO and MMH tanks, since tank pressure drift and LV21/LV31 back-pressure relief were previously covered in Sections VI and VII, respectively.

The most noteworthy propulsion system anomaly during two decades of virtually flawless operation was prime pressure regulator (PR1) leakage early in the mission. PR1 leakage was covered in detail in Ref. 6, but to allow this paper to be a stand-alone document, Table 16 below briefly summarizes the history of Cassini PR1 leakage in flight. Early mission PR1 leak rates were a whopping 3-4 orders of magnitude higher than the leakage spec of 0.6 sccm (standard cubic centimeters per minute), but excellent leak-tight LV10 performance in flight mitigated PR1 leakage concerns greatly. Cassini prime regulator leakage was observed only three times throughout the mission, during initial pressurization and the helium flow events just before and after DSM. LV10 closing timing was carefully selected for all future pressurization events; in fact, it was conservatively picked so the PR1 hard seat would never close again in flight. As such, there was no concern per se for future PR1 leakage after DSM, but this also means PR1 leakage rates were not able to be determined again in the mission, plus LV10 leak tightness then became enormously important.

Table 16 lists the dates of PR1 leakage, helium tank initial and final pressures, PR1 leak rates, and inferred contaminant particle size diameters consistent with observed leak rates. For initial pressurization and the pre-DSM pressurization, there were two ways to calculate PR1 leak rate, from (1) NTO/MMH tank pressure rise rates or (2) upstream helium pressure vs. time decay curves. Amazingly, inferred PR1 leak rates calculated via these two methods agreed with each other within a mere 2%! The bad news, though, was the magnitude of the leakage itself. The PR1 initial leak rate was around 1636 sccm, over 2700 times larger than the leak spec of 0.6 sccm. Despite this grossly out-of-spec performance, a tiny 0.18-μm particle stuck at the PR1 hard seat would be sufficient to cause PR1 leakage at this rate. Such a particle would be nearly 100x smaller than the filter rating of 15 μm, readily passable by

filter F2 (see Figure 8). Minor debris swept downstream following PV1 firing, either miniscule pre-existing contaminants in the lines or PV1 pyro blowby particles, could easily explain PR1 leakage at initial pressurization.

Even with PR1 gross leakage expected to persist around DSM, the flight team was surprised yet again when the inferred leak rate increased by roughly a factor of 6.6 vs. the initial pressurization PR1 leak rate. Actual PR1 leak rates decreased 15-20% pre-DSM to post-DSM, but the inferred particle size was basically the same before and after DSM (given decreases in helium supply pressure across DSM), around 1.2-1.3 μm in diameter. Note this is still more than ten times smaller than the F2 filter rating of 15 μm, so particulate contamination at the PR1 hard seat remains the most likely cause of increased PR1 leakage before and after DSM vs. the initial pressurization leak rate.

Following DSM, the regulator outlet pressure showed an anomalously high value. This was explained as a consequence of closing LV20 and LV30 ten and twenty seconds, respectively, after the closure of LV10. For regulated, main-engine burns longer than sixty minutes, it was desirable to close LV20 and LV30 as soon as possible to avoid possible cold traps for NTO migration and condensation. Both LV20 and LV30 have forward pressure-relief functionality, and this forward-relief feature was inadvertently tested following DSM. In the presence of a leaking regulator, it would be theoretically possible to overpressurize the PR1 outlet if LV20 and LV30 were closed too early and they did not forward relieve. This idiosyncrasy in the face of PR1 leakage was only recognized after the fact, unfortunately.

An independent anomaly in post-DSM data playback precluded the gathering of 1-Hz telemetry samples for PR1 outlet pressure. Regrettably, only real-time values were available, updated just once every sixty-four seconds This exacerbated attempts to determine if the PR1 outlet might have been overpressurized following LV30 closure. Fortunately, the first PR1 outlet pressure telemetry sample was only eight seconds after LV30 closure, and it showed a pressure value exactly consistent with the forward pressure relief of LV20 and/or LV30. Even in the pathological case of having LV20 and LV30 not relieve until precisely that point (eight seconds after LV30 closure), as required by PR1 outlet pressure telemetry, the maximum PR1 outlet pressure would have "only" been 657 psia, slightly below the proof pressure of 685 psia. There is no reasonable expectation that LV20 or LV30 failed to pressure relieve as designed, but this worst-case assessment allayed any fears that the regulator might have been damaged post-DSM.[6] Flawless operation of PR1 at TCM-13 and TCM-20 negated any residual concerns as well.

Table 16: Summary of Cassini Primary Pressure Regulator (PR1) Leakage in Flight

Pressurization Event	**Initial Pressurization**	**Pre-DSM Pressurization**	**Post-DSM Pressurization**
Date [M/D/Y]	11/9/97	12/2/98	12/3/98
Helium Tank Supply Pressure PH1 Δ [psia]	3449 → 3323	3290 → 3231	3231 → 2433
PR1 Leak Rate from P(NTO)/P(MMH) Increase	1636 sccm	11251 sccm	9311 sccm
PR1 Leak Rate from PHE2 Pressure Decrease	1667 sccm	11113 sccm	N/A (LV20/30 closed)
Error between PR1 Leak Calculation Methods	+1.9%	-1.2%	N/A
PR1 Leakage Specification	0.6 sccm	0.6 sccm	0.6 sccm
Observed Leak Rate/Specification Leak Rate	2727x	18752x	15518x
Particle Diameter Consistent w/ Leak Rate [μm]	0.18	1.2	1.3
Maximum Particle Diameter thru Filter [μm]	15	15	15
Particle @ PR1 Hard Seat Could Explain Leak?	Yes	Yes	Yes

It is readily apparent from the CPMS schematic (Figure 8) there is a shared pressurization system for NTO and MMH, a design simplification but also a source of potential concern for a mission as lengthy as Cassini. The CPMS design is a product of its time, and the mid-1990's were notorious for bipropellant pressurization system issues on other planetary missions. The very public loss of Mars Observer in August, 1993 brought vast attention to the potential for NTO migration, possible condensation, and even energetic reaction with MMH in the upstream portions of shared bipropellant pressurization systems. Even though root cause for the Mars Observer failure was never identified, NTO vapor migration issues remain the top candidate.[14] A few years later, Galileo experienced bipropellant pressurization system issues of its own. These included a "sticky" NTO check valve during its second main-engine burn (the Orbiter Deflection Maneuver, ODM) and an MMH check valve flow restriction greater than

50% during Galileo's fourth and final main-engine burn, the PeriJove Raise (PJR) maneuver. The latter was attributed to the formation of Fuel-Oxidizer Reaction Products (FORP) all the way over at the Galileo MMH check valve, due to NTO vapor migration.[3]

Given these snafus with mid-1990's bipropellant pressurization systems on other planetary missions, the CPMS design was modified to include many mitigating features to protect against NTO migration issues upstream of the propellant tanks. These included PCA and helium line heaters (precluding upstream NTO condensation by keeping temperatures warmer than the NTO tank), low-pressure latch valves LV20/LV30, quad-redundant Sterer check valves CV20-23/CV30-33 with fantastic reverse leakage characteristics, and even an NTO-side pyro-isolation ladder, PV22-29. The general strategy was to use PV22-29 to absolutely limit opportunities for NTO migration through CV20-23 and LV20, and this plan worked very well in flight. Incidentally, bipropellant pressurization system anomalies continued to occur on other planetary spacecraft after Cassini launched, including the Japanese mission to Venus, Akatsuki (delaying Venus Orbit Insertion for years) and the Juno mission to Jupiter (delaying mission science goals, given the unacceptable project risk of lowering perijove via main-engine firing in the presence of a sticky check valve).[15]

In general, bipropellant pyro firings PV1 and PV22-PV26 were timed to minimize exposure of the upstream bipropellant plumbing to NTO vapor migration. PV1 was not fired until a few days before TCM-1, but downstream burst disks actually precluded NTO vapor migration during the first few weeks of the mission, anyway. The timing of future pyro ladder firing events was more critical. PV23 and PV24 were fired opened and closed, respectively, over the course of just a few days enveloping DSM. It would be another 5.5 years before PV25 was fired, bringing the pressurization system on-line just before TCM-20. However, there was then some discussion about the need to pyro-isolate the NTO tank between TCM-20 and SOI and/or SOI and PRM. To preserve pyro-ladder redundancy, especially given earlier gross PR1 leakage and uncertain future LV10 sealing performance, it was decided to fire PV26 only after PRM had been completed, rather than to pyro-isolate the NTO tank after TCM-20 and/or SOI.

Following PRM, another systems-level trade was performed regarding the timing of PV26 firing. Propulsion engineers favored isolating the NTO tank as soon as possible, but a pyro firing anomaly could well endanger the unique Huygens mission. Given the many levels of mitigation provided on Cassini against NTO vapor migration, the project decided to delay PV26 firing until just after the successful completion of the Huygens mission. Therefore, the CPMS was vulnerable to upstream NTO vapor migration between TCM-20 and the PV26 firing, a period of 255 days. Of course, TCM-20, SOI, and PRM flowed helium through both NTO and MMH upstream plumbing, so any NTO vapor which had diffused upstream was flushed back into the NTO tank during these burns. Since no more regulated OTMs were planned after PRM, the project accepted the risk of NTO vapor migration between PRM and the PV26 firing, a time period of 165 days, in order to minimize Huygens mission risk. One final pyro firing within the CPMS took place fifteen months later, the firing of PV40 to recharge the hydrazine tank. Table 17 below presents a summary of all CPMS pyro firings during the twenty-year Cassini mission; all pyro firings were completely nominal.

Table 17: Summary of Cassini's Propulsion Module Subsystem (CPMS) Pyro Firing Events

Pyro Device [see Fig. 8]	Pyro Fired Open or Closed	Firing Date [M/D/Y]	Purpose of Pyro Firing	Result
PV1	Open	11/8/97	Initial Biprop Pressurization & Break Burst Disks	Success
PV22	Closed	11/10/97	Isolate NTO Tank from MMH side after TCM-1	Success
PV23	Open	12/2/98	Unisolate NTO Tank for pre-DSM Pressurization	Success
PV24	Closed	12/5/98	Isolate NTO Tank from MMH side after DSM	Success
PV25	Open	5/25/04	Unisolate NTO Tank for TCM-20	Success
PV26	Closed	2/4/05	Isolate NTO from MMH after SOI/PRM/Huygens	Success
PV40	Open	4/10/06	Recharge N_2H_4 Tank for Low Titan Flybys	Success

Even given the incredible lengths to which Cassini was protected against NTO migration upstream of the oxidizer tank, representative NTO diffusion calculations were undertaken during the mission. For simplicity, a classic 1D semi-infinite diffusion analysis was investigated initially; the solution to this very famous diffusion

problem is the familiar complimentary error function, erfc( ). This case is likely not very realistic, given excellent reverse leakage characteristics for Sterer check valves CV20-23, but it is simple and it represents the worst case with respect to the upstream position of the NTO diffusion front within common NTO/MMH upstream plumbing.

Figure 42 displays the fraction of saturated NTO vapor as a function of position upstream of the NTO tank, with ROYGBIV rainbow colors in reverse order representing the evolution in time of the NTO diffusion front. The time-zero condition is saturated NTO vapor at the tank inlet (x = 0 on the plot), with zero NTO vapor for all upstream positions (i.e., for x > 0). This should be the exact initial condition right after a regulated burn, given such burns flush the system with helium from the HTA. Time periods between TCM-20 and SOI (34 days), SOI and PRM (53 days), and PRM and PV26 firing (165 days) are included in Figure 42. Note well the NTO diffusion front definitely arrived at CV20-23 within a few days to weeks, and it was not much worse between PRM and PV26 firing than it was between SOI and PRM.

Given the curves in Figure 42, it is also possible to integrate the NTO density profile over upstream position location for a fixed diffusion time. The result of this integration is the maximum amount of NTO vapor mass passing a given upstream location during that time. Future projects with concerns about NTO vapor migration may benefit from performing these calculations as well. As a real world example, NTO diffusion calculations were also performed for Juno, using only minor modifications to this Cassini 1D semi-infinite model. However, the likely poor assumption that CV20-23 pass NTO vapor unfettered led the Juno project to request a 1D finite model be developed in 2009, one that passed zero NTO vapor upstream of their oxidizer check valves. This was a much more difficult problem to solve closed form, but to the author's delight, an elegant solution was found using superposition of solutions of the diffusion equation, a Fourier series expansion of an initial condition, etc. The final solution was an aesthetically pleasing infinite series which satisfied all initial and boundary conditions. Compared to the 1D semi-infinite model, the new boundary condition for the finite case is similar to an insulated wall in the heat equation—namely, no NTO vapor concentration (density) gradient may exist at CV20-23, if the check valve quad pack does not pass any NTO vapor upstream.

Figure 43 below is the finite model analogue to Figure 42. Even though this model was not developed until 2009, it was applied retroactively to the Cassini bipropellant pressurization system for the period between TCM-20 and PV26 firing back in 2004-2005. As expected, the oxidizer diffusion front arrives more quickly to CV20-23 than for the semi-infinite model; essentially, NTO has "nowhere else to go" if CV20-23 stop further NTO vapor migration (cf. Figure 43 to Figure 42). Also as expected, the diffusion curves in Figure 43 all have zero slope at the CV20-23 position location. Finally, note well NTO vapor was nearly saturated at CV20-23 by the time of PV26 firing, but it was more than 50% saturated between SOI and PRM. Put another way, whatever risk was already accepted by not pyro-isolating the NTO tank between TCM-20 and SOI or SOI and PRM was not made much worse by waiting to fire PV26 until after the successful completion of the Huygens mission.

As before, it is possible to integrate the density profiles in Figure 43 over upstream position for a fixed diffusion time. Interestingly, even though the 1D semi-infinite model is conservative with respect to the location of the NTO diffusion front, it is not conservative with respect to accumulated quantity of NTO vapor at a given upstream location. This is as expected, given the inability for NTO vapor to "escape" further upstream. Of course, the most important consideration is to prevent NTO from intermingling with MMH in the first place, so the conservative and multi-pronged CPMS vapor migration design was likely the most important factor in assuring bipropellant pressurization system safety over this lengthy mission.

Incidentally, a version of the 1D finite diffusion model could also be applied retroactively to N_2H_4 vapor diffusion upstream of the hydrazine tank between launch and recharge, helping to determine the maximum amount of condensed hydrazine that was near PV40 before it was fired in 2006. Unlike the bipropellant pressurization system plumbing, temperatures upstream of the hydrazine tank were colder than liquid monopropellant, so N_2H_4 condensation at PV40 was possible. Fortunately, using bracketing conditions from the 1D semi-infinite model, which was available in 2006, the maximum amount of condensed N_2H_4 at PV40 just before pyro firing was estimated to be only around 2-16 mg, alleviating firing concerns related to liquid hydrazine shock detonation (given such a small quantity of N_2H_4). This rather wide range in accumulated N_2H_4 mass of 2-16 mg comes from integrating 8.5 years of hydrazine diffusion from the tank to PV40 or to an infinite upstream position, respectively. These cases should bound the value which would have been obtained from a finite model. However, all of this was of little comfort on April 10, 2006, when PV40 was actually fired. An unrelated but incredibly ill-timed Deep Space

Network (DSN) ground station outage, precisely at the moment of pyro firing, emulated loss of the spacecraft due to PV40 firing! Thankfully, data flow was restored four minutes later, but these few minutes aged the flight team many years.

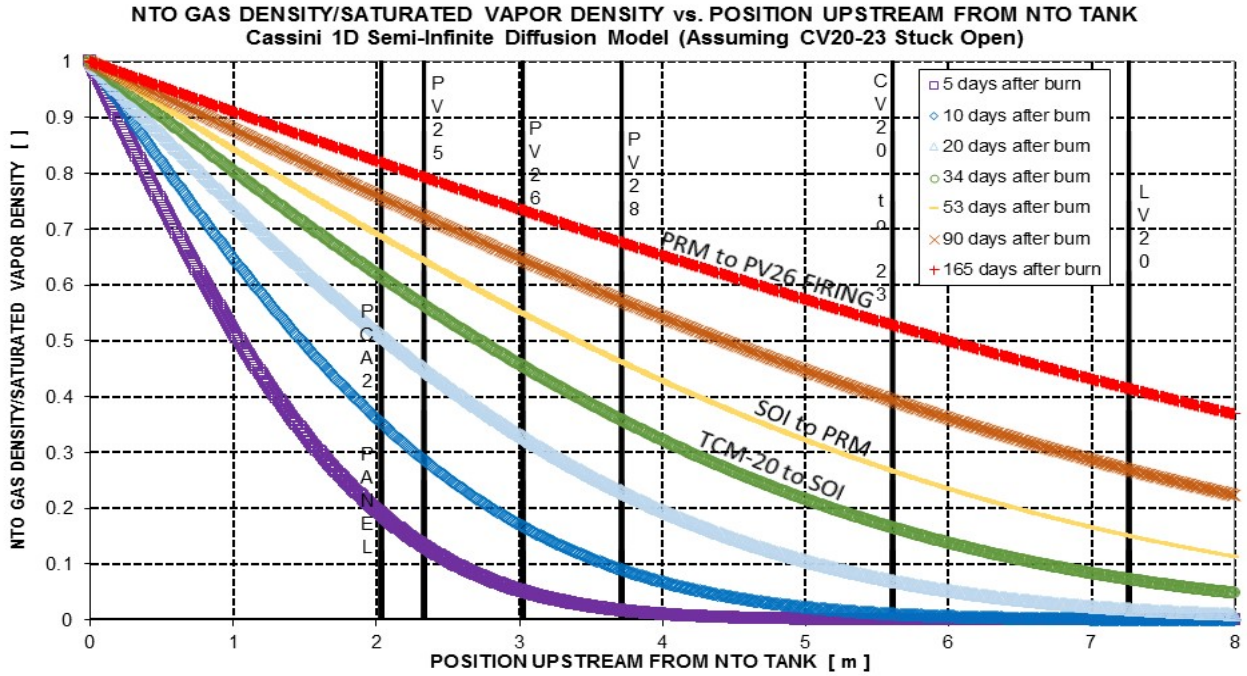

Figure 42: Fraction of Saturated NTO Vapor Density vs. Position Upstream of NTO Tank (Semi-Infinite Case)

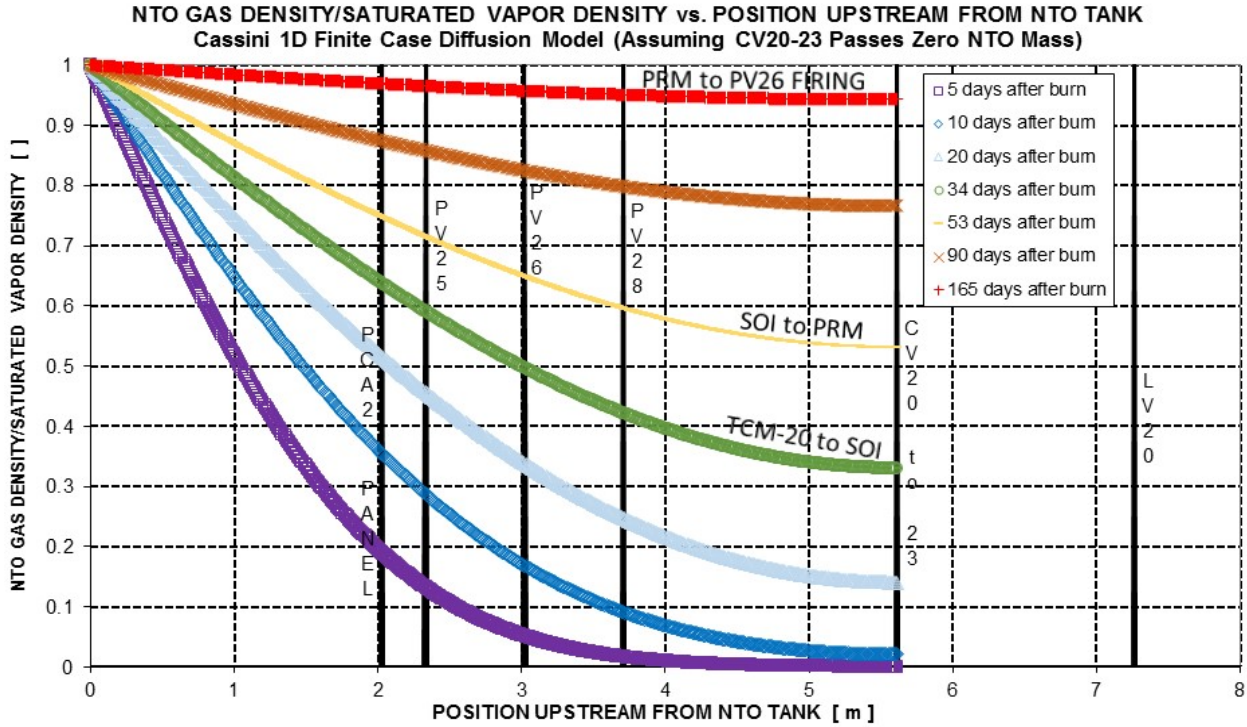

Figure 43: Fraction of Saturated NTO Vapor Density vs. Position Upstream of NTO Tank ("Perfect" CV20-23 Case)

Turning now to bipropellant pressurization system hardware further upstream from CV20-23, LV20 and LV30 offered additional protection against NTO vapor migration over the Cassini mission. As richly instrumented as Cassini was, regrettably there were no pressure measurements between LV20 and CV20-23 or between LV30 and CV30-33. With LV20/LV30 closed for nearly the entire mission, this clouded attempts to assess CV20-23 and CV30-33 opening ("cracking") and closing ("reseat") behavior. However, flight data suggest LV30 (and probably LV20) were leaky enough such that pressures often equilibrated across the low-pressure latch valves, allowing a direct comparison of PREG and MMH tank pressure to determine CV30-33 cracking and reseat performance. This was determined by investigating plots of inferred helium mass between LV10 and CV20-23/CV30-33 as a function of time. This small volume (228.7 cc total, with 162.4 cc of this between LV10 and LV20/LV30) ended up being an incredibly sensitive means by which to monitor potential LV10 helium leakage or diffusion, LV30 leakage, and CV30-33 cracking and reseat behavior over twenty years.

Calculating inferred helium mass in the volume between LV10 and CV20-23/CV30-33 was straightforward, given telemetered values of PR1 outlet pressure (PREG, not available on Galileo) and multiple temperature transducers for the lines. In fact, a very good representative temperature for this volume was able to be determined by averaging one PR1 temperature and three PCA panel temperature measurements. Again using the Beattie-Bridgman equation of state, helium mass values were plotted over the entire mission. With no helium flow through LV10, LV20/30, and CV20-23/CV30-33, helium mass in this volume should have remained constant. Deviating trends in inferred helium mass during quiet periods suggested possible LV10 helium diffusion, LV30 leakage, and/or CV30-33 "crackings" (openings) or "reseats" (closings), depending on the relative leak rates of LV30 vs. LV10. Note that it was difficult to absolutely quantify LV20 or LV30 leakage, given a lack of pressure knowledge between these latch valves and CV quad packs. Moreover, the NTO tank pyro-isolation during most of the mission hindered in-flight assessment of LV20 leakage.

Even with LV20/30 closed and the NTO tank pyro-isolated, there were times when the inferred helium mass in the volume between LV10 and CV20-23/CV30-33 showed nearly step-function decreases. The only logical explanation for this behavior was the opening or cracking of CV30-33, followed by helium flow into the MMH tank. The CV30-33 quad pack cracking spec is 14.0 ± 2.0 psid, and in-flight data were generally consistent with ground test. Figure 44 is a bar chart showing the thirteen suspected CV30-33 crackings during the Cassini mission, with most crackings occurring fairly close to spec. CV30-33 crackings were most often caused by MMH tank cooling, large blowdown maneuvers, or LV30 openings later in the mission, when LV30 was apparently less leaky than earlier epochs and thus held pressure.

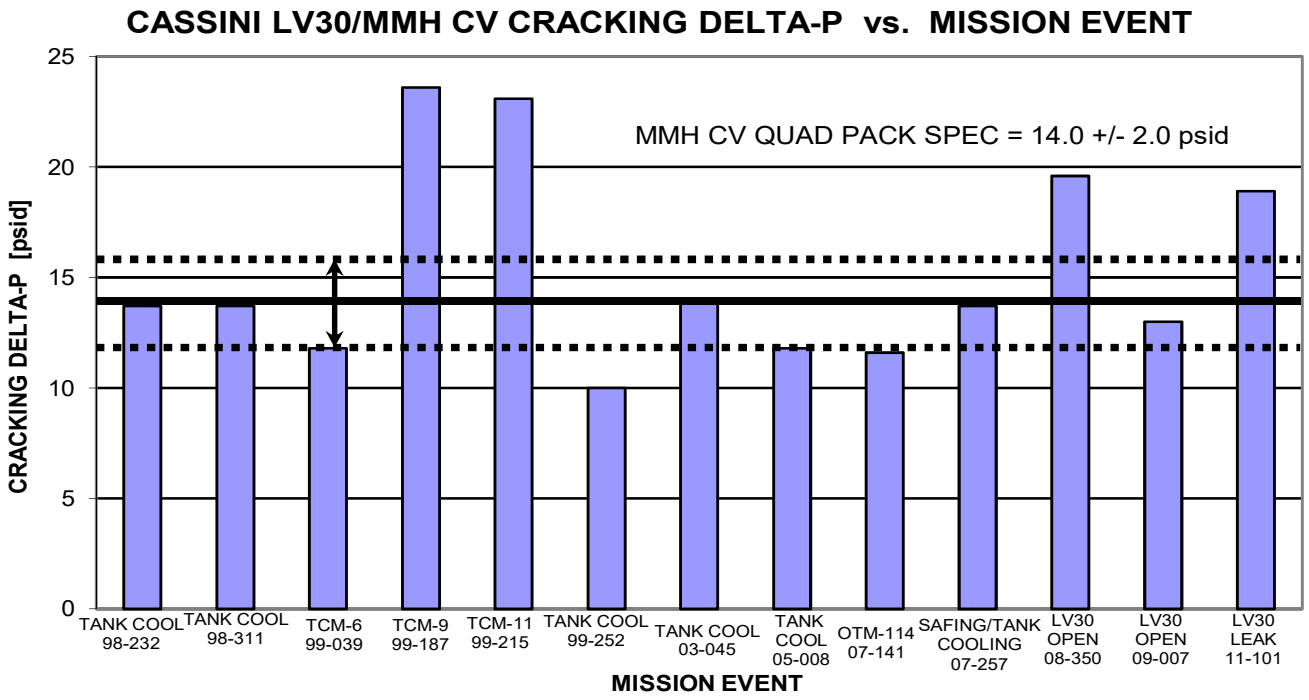

Figure 44: CV30-33 Cracking ΔP vs. Mission Event

The first two CV30-33 crackings in Figure 44 occurred right in the middle of the spec range, and they were caused by MMH tank cooling. For a fixed upstream pressure, cooling the MMH tank decreased fuel pressure, which then placed a larger Δp across CV30-33, thus leading to cracking. Future tank cooling events which caused fuel check valve quad pack openings were similar. In contrast, another way to crack CV30-33 was to initiate MMH tank pressure decreases by executing large blowdown main-engine burns. TCM-6, TCM-9, and TCM-11 all caused CV30-33 cracking, as may be seen in Figure 44. In fact, MMH tank pressure decreases were relatively rapid during the last two of these TCMs, which were quite large blowdown burns (43.5 and 36.3 m/s, respectively). This probably explains why cracking Δp was 7-12 psid above spec for these events, given CV30-33 cracking was very likely a gradual process.

In preparation for the second and final fuel-side repressurization of tour on January 8, 2009, LV30 was cycled open and closed a few weeks in advance to test out long-dormant hardware. Since prior CV30-33 cracking events all occurred with LV30 closed, even though PR1 outlet pressure was 19 psid above MMH tank pressure at that time, propulsion engineers did not expect CV30-33 to crack when LV30 was opened. A strong CV30-33 cracking did transpire during this test, with weaker MMH quad check valve cracking when LV30 was opened a few weeks later in preparation for actual fuel-side repressurization. This again highlights how useful it would have been to have pressure telemetry between LV30 and CV30-33. It is also worth noting DN resolution for MMH tank pressure and PR1 outlet pressure transducers was only 1.88 psia/DN, so error bars on the data of Figure 44 would be fairly sizable.

One final CV30-33 cracking took place in 2011, with LV30 closed and no MMH tank or PR1 outlet pressure changes. This "out-of-the-blue" event was initially quite puzzling, but a reasonable hypothesis was eventually cobbled together. Perhaps minute LV30 leakage after the final fuel-side repressurization eventually cracked CV30-33 (given constant MMH tank pressure), causing the pressure between LV30 and CV30-33 to decrease. This would have placed more Δp across LV30, increasing its leak rate (presumably), eventually causing PR1 outlet pressure to drop, which then showed up in the helium mass calculations for this volume. It is a convoluted picture, but it is consistent with the data and with nominal performance for all CPMS upstream components. In summary, CV30-33 cracking performance was excellent throughout the Cassini mission, likely due in no small part to extensive efforts to mitigate NTO migration.

After CV30-33 crackings occurred, further Δp decreases across the quad pack were required to close or reseat CV30-33. Within plots of line helium mass vs. time, reseats were generally marked by helium masses which "flattened out" after the aforementioned step function decreases. Quite often, reseat was a much more gradual process than cracking. One reason for this is the first ten of thirteen CV30-33 reseats in flight were caused by tank heating, which slowly raised MMH tank pressure and thus incrementally lowered Δp across CV30-33 for fixed upstream pressure. Figure 45 is the reseat analogue to Figure 44, and it includes the CV30-33 closing spec of 12.0 ± 2.0 psid. In actuality, the reseat spec is the cracking spec minus 2 psid, so if cracking Δp in flight were higher or lower than spec, reseat Δp would be expected to follow suit.

In general, CV30-33 reseat occurred below spec during Cassini's lengthy mission, as may be seen from Figure 45. This is the direction of "sluggish" closing response of the check valve quad pack, but this had no implications for the mission. The final three reseats in Figure 45 were actually caused by helium flow after CV30-33 cracking, rather than by MMH tank heating. Since helium flow caused more rapid pressure changes than were caused by tank heating, these reseats seemed to be stronger than the first ten reseats. As such, perhaps it is not surprising these reseat Δp's were closer to specification. In conclusion, CV30-33 reseat behavior in flight was nominal, if perhaps occurring a bit below specification levels for "leisurely" reseats. Though not covered heretofore, flowing Δp values for CV30-33 (and CV20-23) during regulated burns and fuel-side pressurizations were also nominal. These Sterer check valve quad packs displayed stellar performance over decades, and their flawless operation helped Cassini execute its mission impeccably.

To aid future missions, it might be useful to archive twenty years of Δp data between the PR1 outlet and the MMH tank. Figure 46 below charts this difference [PREG – P(MMH)] as a function of mission time. For clarity, only the CV30-33 cracking spec is included in Figure 46; recall the CV30-33 reseat spec is just 2 psid lower than the cracking spec. Suspected cracking events are notated in Figure 46 within ellipses, while inferred reseats are marked by rectangles. There is a one-to-one correspondence between the thirteen cracking and reseat events in Figures 44-45 with the ellipses and rectangles in Figure 46. One stark change in Δp not previously discussed is readily apparent in Figure 46—a rapid decrease from 25 to 16 psid around mission day 6077. This was caused by turning off the PCA

panel and helium line heaters, which lowered upstream temperatures (and thus PREG), all while keeping MMH tank temperature and pressure basically constant. Despite the lack of pressure telemetry between LV30 and CV30-33, it is noteworthy how much insight into CV30-33 performance was possible via simple gas law calculations and a modicum of telemetry channels.

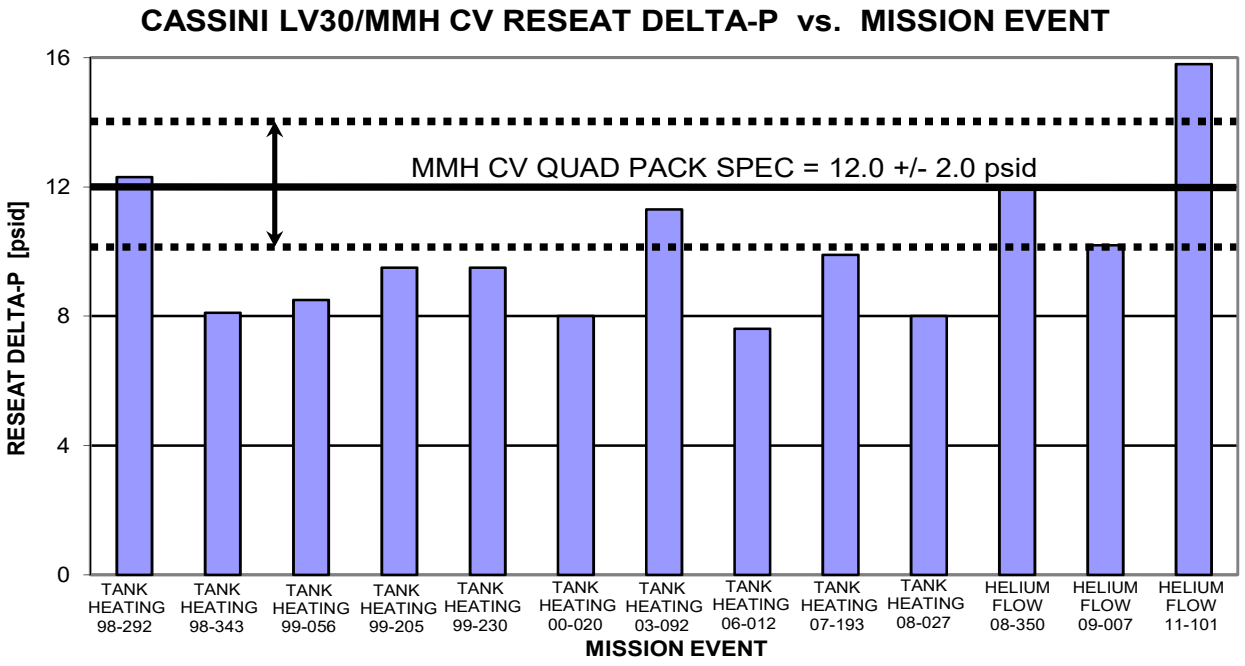

Figure 45: CV30-33 Reseat ΔP vs. Mission Event

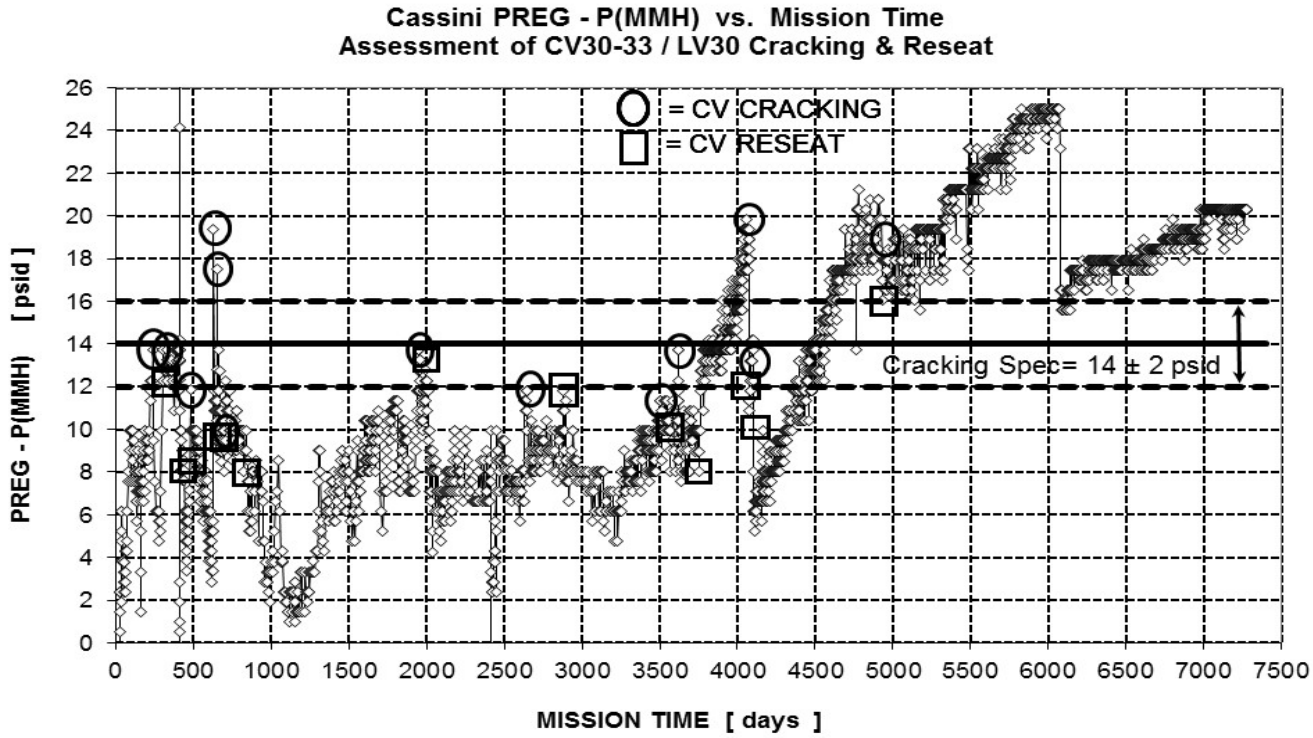

Figure 46: PREG – P(MMH) vs. Mission Time

Amazingly, the inferred helium mass data between LV10 and CV20-33/CV30-33 can be further "tortured" to potentially assess LV30 leakage itself. Specifically, any previously identified cracking event found from status report data (with a frequency typically no better than once per day) was thoroughly investigated, using all available telemetry samples from PR1 outlet pressure, PCA panel temperatures, and PR1 temperature around the event. Most of the time, decaying exponential curve fits of helium mass vs. time matched data points quite well, with the initial slope of the exponential fit a reasonably accurate characterization of worst-case LV30 forward leakage rate. The LV30 leak spec was 5 scch, and there were a few instances in flight where LV30 apparently leaked above spec, though not grossly.

The entire history of Cassini LV30 leakage, including ground-test data, is presented below in Figure 47. This bar chart displays leak rates on a semi-log scale, given the wide range in measured or calculated leakage over LV30's long history. Note well there were no spec violations before launch, and in-flight performance was generally quite similar to ground-test performance. The three events in flight during which LV30 may have leaked higher than specification were TCM-9 ("Crack #4"), TCM-11 ("Crack #5"), and after LV30 was cycled in late 2008 in preparation for the final fuel-side repressurization of tour ("Crack #11"). All three of these events had high CV30-33 cracking Δp, as may be verified in Figure 44, so perhaps it is not surprising LV30 leaked more helium, given larger pressure differentials across the latch valve, presumably.

Regrettably, LV20 and CV20-23 performance in flight was rarely observable, given the NTO tank was pyro-isolated from upstream components for the vast majority of the Cassini mission. In theory, CV20-23 crackings and/or reseats between TCM-20 and the PV26 firing should have been visible. However, PREG – P(NTO) values during this time were quite low, never exceeding 7 psid during quiescent periods. As such, CV20-23 was likely closed throughout this nine-month timeframe, other than during regulated burns TCM-20, SOI, and PRM. Therefore, LV20 ground-test data alone will have to suffice for assessing leakage. Figure 48 below is the LV20 analogue to Figure 47, on the same vertical scale. Note well there were some very minor LV20 leakage spec violations before launch, unlike LV30, but these had no consequences for mission operations. Figure 48 is included for completeness only.

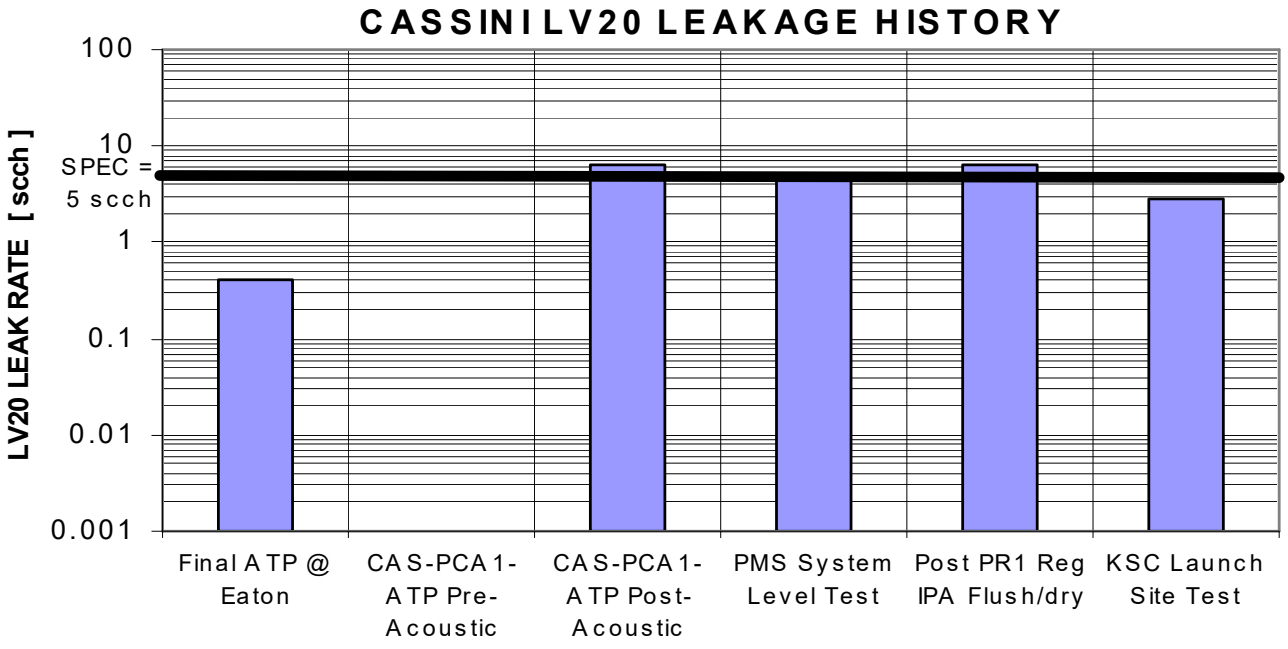

Figure 47: Cassini LV30 Leakage History

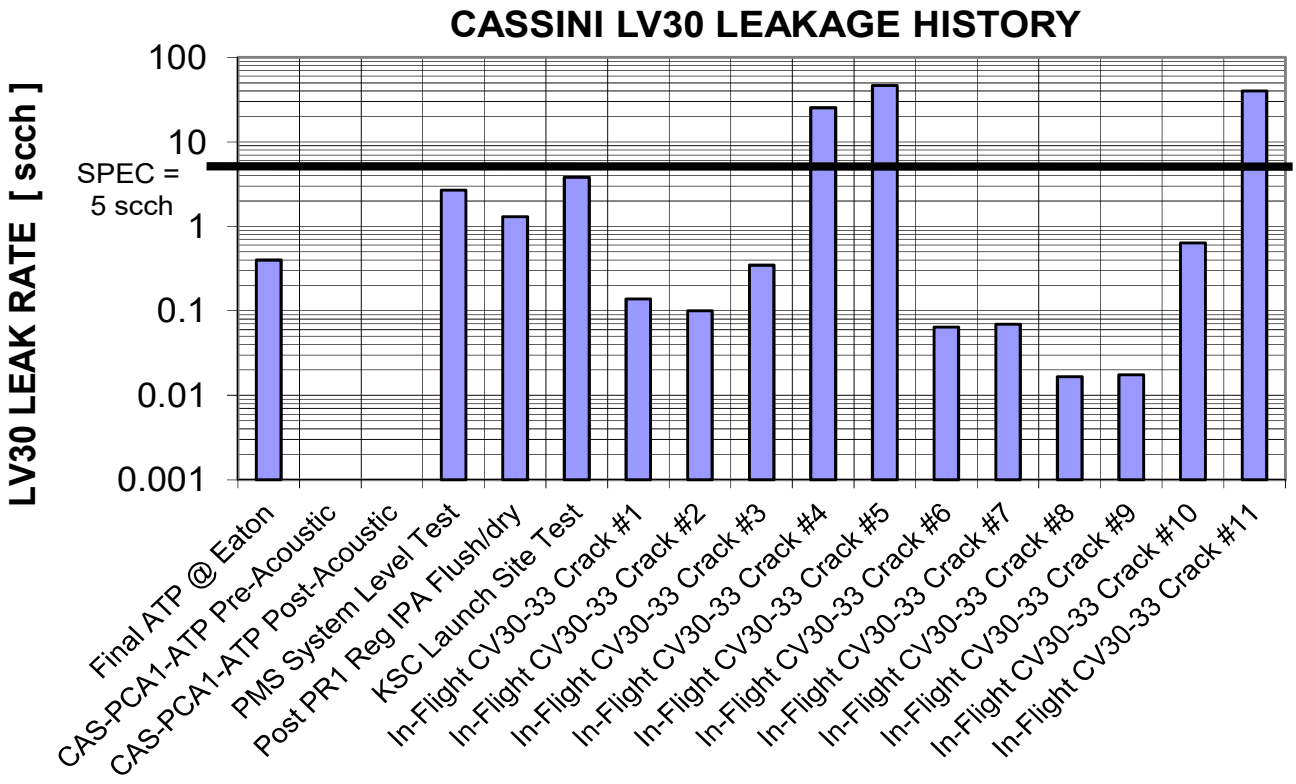

Figure 48: Cassini LV20 Leakage History (Pre-Launch Data Only)

Another viable way to present the LV30 leakage data from Figure 47 is in a tabular format. One reason to do this is to capture the in-flight time frames and durations over which LV30 leakage took place, variables not able to be shown in Figure 47. Table 18 summarizes post-launch LV30 leakage rates over the mission, again punctuated by three minor spec violations over two decades. Since there were no post-launch data points available for LV20 leakage, there is no need to similarly tabulate the data from Figure 48.

Table 18: Cassini LV30 In-Flight Leakage History

Dates of Leakage	Mission Time [days]	Maximum Leak Rate [scch]	Spec Leak Rate [scch]	Meets Spec?
98-231 to 98-253	308 to 330	0.138	5.0	Y
98-311 to 98-330	388 to 407	0.101	5.0	Y
99-035 to 99-049	477 to 491	0.347	5.0	Y
99-187 to 99-188	629 to 630	**25.5**	5.0	**N**
99-214 to 99-215	657 to 658	**46.5**	5.0	**N**
99-243 to 00-017	685 to 824	0.0639	5.0	Y
03-045 to 03-092	1948 to 1995	0.0694	5.0	Y
05-008 to 06-012	2644 to 3012	0.0166	5.0	Y
07-141 to 07-193	3506 to 3558	0.0174	5.0	Y
07-257 to 08-027	3622 to 3757	0.635	5.0	Y
11-101 to 11-102	4926 to 4927	**40.2**	5.0	**N**

Returning to LV30, rather than trying to cross-correlate the data provided in Figures 46 and 47, it might be more enlightening to plot LV30 inferred leak rate directly as a function of [PREG – P(MMH)]; this is done below in Figure 49. As expected, LV30 leak rate appeared to increase as the Δp across the latch valve increased. The forward pressure-relief spec for LV30 was 100 ± 20 psid. Coupled with the CV30-33 cracking spec of 14.0 ± 2.0 psid, this means LV30 should not forward-relieve until [PREG – P(MMH)] reached 114 ± 22 psid. LV30 apparently passed some helium with much less Δp across it, but this was likely due to very minor LV30 leakage rather than premature LV30 forward pressure-relief anomalies. This slight LV30 leakage in flight is also in family with LV20 minor spec leakage violations during ground testing. In summary, even though there were perhaps minor LV30 leakage spec breaches during flight, both LV20 and LV30 performed very well during the protracted Cassini mission.

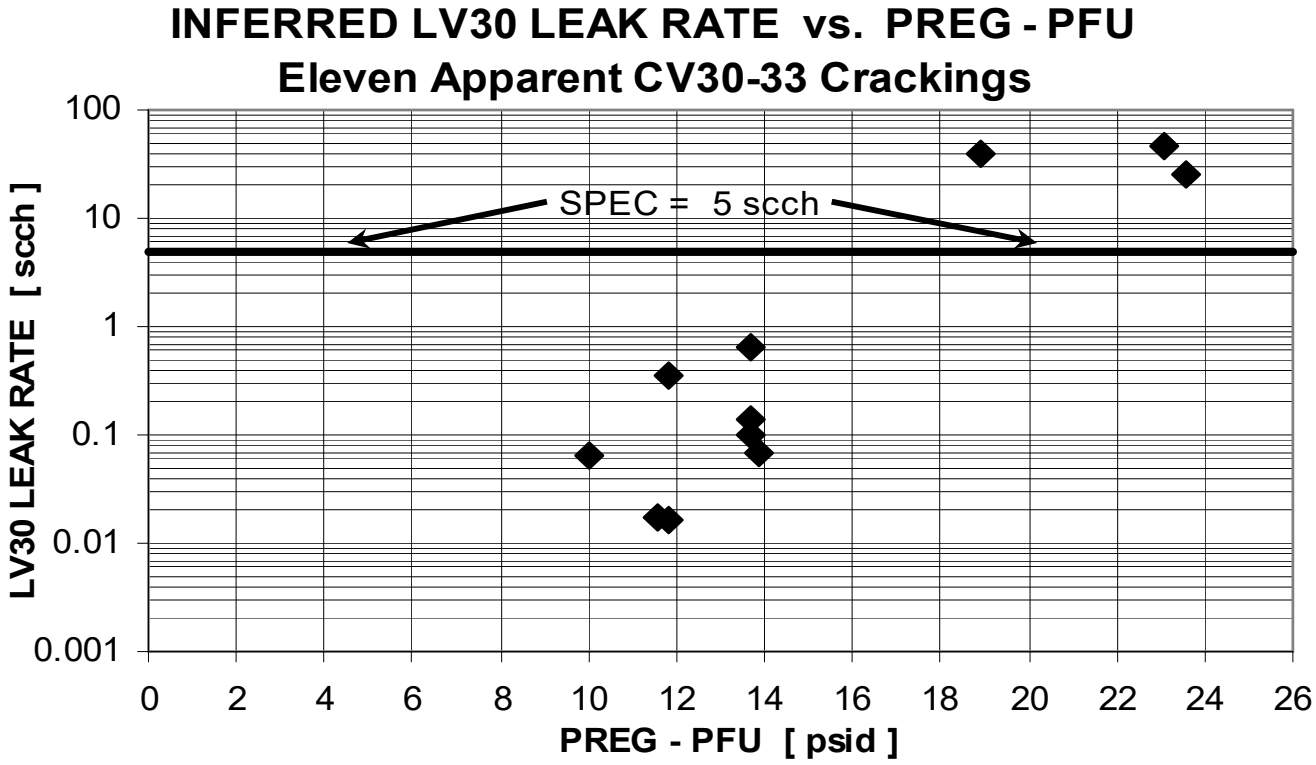

Figure 49: Inferred LV30 Leak Rate vs. PREG – PFU

One assumption was tacitly made throughout this entire section—namely, LV10 leakage was much smaller than LV20/LV30 leakage or CV20-23/CV30-33 leakage. A priori, this might not be expected to be the case, since LV10 was a high-pressure latch valve and thus might be more prone to leakage than low-pressure latch valves or check valve quad packs. Fortunately, LV10 sealing performance in flight was truly exemplary. The small line volume between LV10 and CV20-23/CV30-33 turned out to be an incredibly sensitive LV10 leak detector, even given the somewhat coarse PREG telemetry resolution limit of 1.88 psia/DN. In the 5.5 years between DSM and TCM-20, a subtle but indisputable rise in PREG of about 1 DN/year was noticed. In fact, this observation is what triggered the thought in the first place of trending helium mass in this volume. CV30-33 crackings caused downward step-functions, as was mentioned previously, but during quiescent periods, helium mass actually appeared to ramp up linearly, very repeatably, and with a signal well above noise levels, at least over long mission times. Perhaps LV10 was leaking slightly after all?

Spec leakage for LV10 was 20 scch, four times higher than the LV20/LV30 spec, given it is more difficult to contain high-pressure helium than low-pressure helium. Despite this increased allowable leakage, LV10 sealing performance was astonishingly good in flight. The worst-case measured "leakage" occurred after the LV10 closing

following initial pressurization, and it was at a rate of only 0.00460 scch. This is 4350 times smaller than the leakage spec of 20 scch! Future incarnations of LV10 "leakage" measured in flight were similar but even smaller. In fact, "leakage" rates were so low that helium diffusion through LV10 rather than leakage was suspected. Rough calculations suggested Reynolds and Knudsen numbers for these vanishingly small helium flow conditions were indeed close to 1, confirming operation in the diffusion regime. With decreasing helium tank supply pressure, choked flow considerations across LV10 would dictate that any leak rate would decrease in proportion to the square root of the upstream pressure for a fixed leak area. In contrast, LV10 helium flow rates while closed seemed agnostic to upstream pressure conditions over the mission, another clue that simple helium diffusion through LV10 was responsible for the trends noted. Figure 50 below is the LV10 analogue to Figures 47-48, also including ground-test data. Note well LV10 was apparently leak free during all of ground testing, though the level of leakage detectable pre-launch is unknown. It is doubtful ground tests were of sufficient duration to discern the subtleties of helium diffusion through LV10, the theoretical lower limit for helium flow rate through this closed high-pressure latch valve.

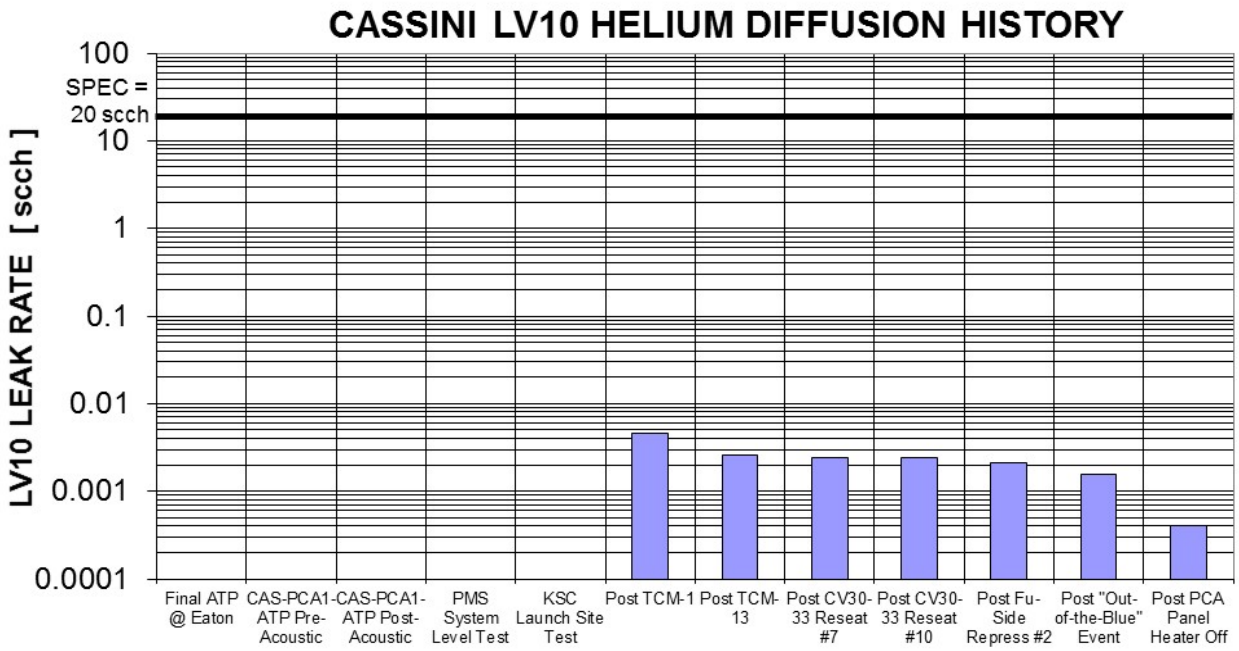

Figure 50: Cassini LV10 Helium Diffusion History

One of the most satisfying observations from Figure 50 is the invariance of the measured helium diffusion rate through LV10, particularly between TCM-13 (August, 1999) and the second fuel-side repressurization of tour (January, 2009). Over this 9.3-year period, upstream pressure decreased markedly from 2433 psia to 887 psia, TCM-20 and SOI and PRM cycled LV10 (along with the first fuel-side repressurization of tour), and CV30-33 cracked and reseated seven times (see Figures 44-45). Despite this plethora of changes in upstream conditions and to the mass of helium trapped between LV10 and CV20-23/CV30-33, inferred LV10 diffusion rates over four measurement cycles during this time period agreed within 20%. This may be more easily seen in Table 19, an equivalent way to present the data from Figure 50. Rows two through five of Table 19 show remarkable agreement in LV10 diffusion rate.

Between TCM-13 and TCM-20, an MMH tank cooling event early in 2003 cracked CV30-33 and caused a step-function decrease in helium mass. Despite this large perturbation to the helium mass in the trapped volume downstream of PR1, the helium mass trend following CV30-33 reseat "picked up where it left off" and continued increasing. LV10 diffusion rates across this cracking/reseat event only changed from 0.00256 scch to 0.00243 scch, a difference of only 5%. Even more amazingly, the change in inferred LV10 helium diffusion rate between 2003-2004 and throughout most of 2008 was only 1%! Upstream conditions were far different between 2003-2004 and

2008, so these data again strongly suggest LV10 leakage was non-existent, with this critical high-pressure latch valve demonstrating diffusion-limited performance during the lengthy Cassini mission. In the presence of large and persistent (and presumably incurable) PR1 leakage, this was most fortunate.

Table 19: Cassini LV10 In-Flight Helium Diffusion History

Dates of LV10 Diffusion [YR-DOY]	Mission Time [days]	Maximum Diffusion Rate [scch]	Spec Leak Rate [scch]	Meets Leak Spec?
97-315 to 98-232	27 to 309	0.00460	20.0	Y
00-005 to 03-045	812 to 1948	0.00256	20.0	Y
03-092 to 04-183	1995 to 2451	0.00243	20.0	Y
08-027 to 08-350	3757 to 4080	0.00245	20.0	Y
09-009 to 11-101	4105 to 4926	0.00213	20.0	Y
11-102 to 14-154	4927 to 6076	0.00158	20.0	Y
14-157 to 17-258	6079 to 7275	0.00040	20.0	Y

Even though LV10 was only opened a handful of times over two decades, and for quite limited durations at that (given PR1 leakage concerns), these LV10 openings enabled one final opportunity for spot-checking Cassini sensors for pressure transducer drift. Specifically, given minimal line pressure drops for low-rate helium flow, HTA tank pressure (PHE1) and helium line pressure (PHE2) should be virtually identical while LV10 was open (see Figure 8). Since PR1 was either leaking grossly or open during the vast majority of the mission, PHE2 agreed well with PR1 outlet pressure throughout the mission while LV10 was closed, as expected. However, the few epochs during which LV10 was open offered a chance to investigate any relative pressure transducer drift between PHE1 and PHE2.

Figure 51 displays the difference between PHE1 and PHE2 as a function of mission time, but with data points only provided during pressurization events, within windows when LV10 was open. A consistent Δp value of 7-8 psid was noted across the mission, with an essentially flat trend, particularly compared to the PHE1 and PHE2 coarse data resolution of 16.8 psia/DN. The ten points in Figure 51 represent the nine separate bipropellant pressurization events during the Cassini mission, with two points at DSM (pre-DSM pressurization and DSM itself). Over more than eleven years, there was no discernible relative drift between PHE1 and PHE2—one last piece of in-flight evidence that Cassini pressure transducers were truly exemplary with respect to sensor drift.

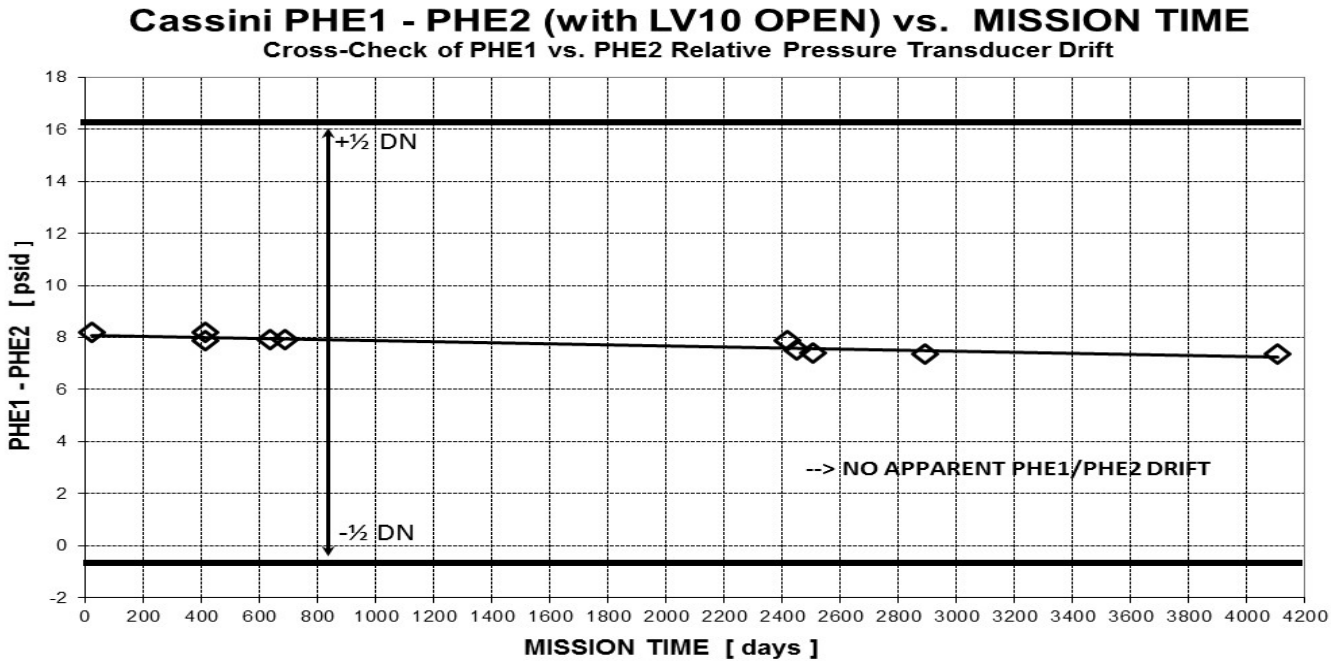

Figure 51: Cassini PHE1 – PHE2 (with LV10 Open) vs. Mission Time

For completeness, a bar chart of PR1's entire leakage history, including ground test data, is presented below as Figure 52. Note well there were minor leakage spec violations pre-launch, on three different occasions (in CPMS system level testing, after flushing PR1 with isopropyl alcohol, and even after IPA drying). Thankfully, the final PR1 leakage measurements at the launch site (KSC) showed zero leakage. In-flight PR1 leakage was covered previously, but the entire leakage history for PR1 is included in Figure 52 for comprehensiveness.

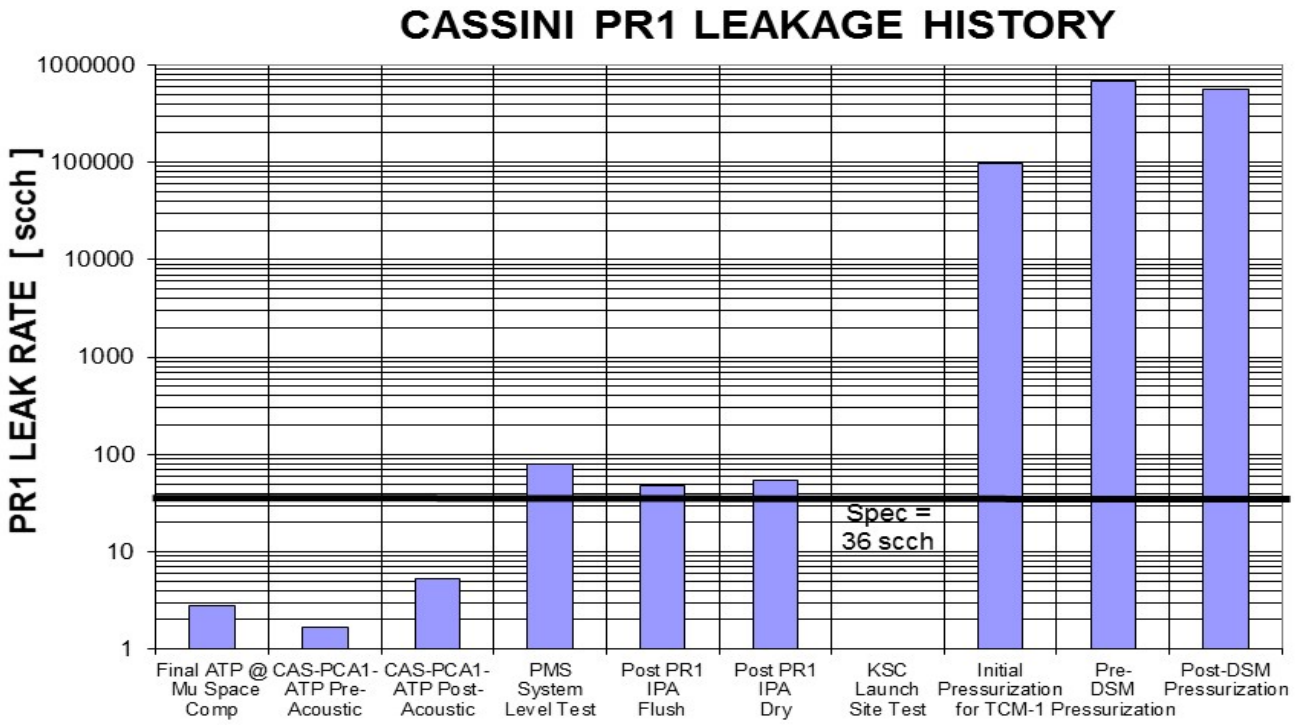

Figure 52: Cassini PR1 Leakage History

Comparing Figure 52 with Figure 50, one general observation was that LV10 leakage performance was about four orders of magnitude better than spec, while PR1 leak tightness was roughly four orders of magnitude worse than spec. As the propulsion team was fond of joking, on a log scale PR1 and LV10 showed nominal average performance over the mission. Two lessons learned seem most pertinent here; namely, (1) hard-seat regulators are prone to leakage, no matter the precautions taken with system cleanliness and (2) high-pressure latch valves offer an excellent back-up sealing capability for leaking hard-seat regulators. Galileo's soft-seat regulator from Mu worked beautifully over many years, but the effort to duplicate Mu's soft-seat regulator for Cassini was unsuccessful, due to many factors (retired and deceased Mu engineering personnel, the lack of detailed as-built drawings, material issues, etc.). Many kudos are due the propulsion teams at JPL and Lockheed Martin Astronautics (LMA) for a truly robust CPMS design, especially the inclusion of high-pressure latch valves LV10 and LV11 (the latter a pristine, unused back-up for LV10).

Despite bipropellant pressurization system performance that was virtually flawless, except for PR1 leakage, external reviewers expressed concerns about PR1's ability to regulate during the critical SOI burn. During the lengthy cruise between Jupiter and Saturn, there were years to reexamine prior decisions and entertain many failure modes. Even though largely thought to be non-credible, a wide-open regulator at SOI surely would end the mission. The project opted to develop a new fault protection algorithm that could allow orbit capture at Saturn, amazingly even in the face of a wide-open PR1 during SOI. The effort required to develop a so-called "bang" controller algorithm was significant, particularly for fault protection, systems, test, and propulsion engineers on the flight team.[16] Thankfully, PR1 regulated perfectly during SOI, precluding the need to invoke this last-ditch effort to ensure Cassini would become the first artificial satellite of the ringed planet.

To very broadly summarize this section, despite the inherent complexity of the CPMS design, particularly in the upstream portion of the bipropellant plumbing, this great effort seems to have paid off handsomely. There are precious few propulsion systems which have operated as immaculately as the CPMS, let alone for two decades in flight. The success of the CPMS, and the Cassini mission itself, may well be a direct consequence of the robustness of the CPMS design, particularly given a myriad of propulsion system issues that have arisen on other missions.

XI. Cassini Propulsion Module Subsystem Tank Pressure and Temperature History

The results from the last three sections are more fundamental than simple time history plots of telemetered tank pressure and temperature, in the author's opinion. One reason for this is tank pressure changes follow tank temperature changes very closely, in a predictable way. Derived quantities, such as propellant or helium mass, should offer more insight into propulsion system performance, trending, leakage, and health and safety. They also "smooth out" relatively unimportant variations in tank pressure caused by innumerable tank temperature changes over the mission. However, for archival purposes, it is likely worthwhile to capture plots of CPMS tank pressures and temperatures over this lengthy mission. In particular, future missions may well need to understand as-flown pressure and temperature ranges for the CPMS tanks, so these data should be captured for posterity.

In-flight temperature control for the CPMS entailed unique challenges, due to fairly stringent component temperature limits, large thermal gradients, and widely disparate solar distances between 0.7 and 9.6 AU. Despite these challenges, CPMS temperature control during the mission was excellent, as may be verified in the plots below. Thermal control was accomplished using thermal blankets, louvers, permanently powered heaters, variable and fixed radioisotope heater units (RHUs), and Automatic Thermal Control (ATC) algorithms which cycled heaters to control temperatures within narrow temperature ranges.

Compared with other extended planetary missions, power margin issues were truly minimal on Cassini, even after twenty years of Pu-238 decay within the three RTGs. As such, the only permanent CPMS heater modification required during the mission was a turn-off of PCA panel and helium line heaters in 2014, saving roughly 10 W of power. Since the final pressurized maneuver occurred nearly a decade earlier, there was no concern for allowing this portion of the bipropellant pressurization system to cool for good, particularly since PV26 was fired closed in 2005. One final CPMS thermal change was required for power margin purposes during the last few months of the mission, but only three times and only for a few hours each time. There was one spacecraft science operating mode which required copious power, so RCS B-branch catbed heaters were turned off during three small science windows to keep power margin comfortably positive. These brief periods with no catbed heating during the Grand Finale had no consequences for Cassini, since the spacecraft was exclusively under RWA control during this science mode.

Turning now to plots of CPMS tank temperature and pressure over the entire mission, Figure 53 below displays the time history of the recharge tank temperature. Some basic trends visible in Figure 53 apply to all tanks, including (1) a general decrease in temperature vs. time due to increasing solar distance and RTG decay, (2) upward spikes around two perihelia near the Venus-1 and Venus-2 flybys, (3) increased temperatures while under RWA control vs. RCS control (due to central body heating), (4) temperature decreases around safing events, (5) increases in temperature while the main-engine cover was closed (allowing less heat to be radiated to space), and (6) large temperature spikes in both directions from the Propellant Gauging Test (PGT) in 2012. Other notable temperature changes in Figure 53 were caused by periods of spacecraft quiescence (i.e., no orbiter science) around Huygens probe relay and the RCS B-branch thruster swap. Post-launch heating caused a rather large but persistent temperature change very early in the mission. Finally, recharge itself caused a transient decrease in recharge tank temperature, given nearly adiabatic expansion of helium after firing PV40. All data in Figure 53 are as expected, covering an in-flight range of only 18-32°C, well within the non-operating allowable flight temperature (AFT) limits of 5-45°C and comfortably within the operating (i.e., at-ignition or firing) AFT limits of 10-35°C.

Recharge tank pressure is presented below in two plots, one before and one after recharge on April 10, 2006. Figure 54 shows the recharge tank pressure between launch and the moment just before PV40 firing. As expected, the curve closely follows Figure 53. Since Figure 29 showed constant helium mass in the recharge tank for the first 8.5 years of the mission (as expected, given no pressure transducer drift), Figures 53-54 should follow each other

very closely. Many of the same transients noted in Figure 53 are discernible in Figure 54 as well, muted somewhat given the coarse resolution of the RTA pressure transducer.

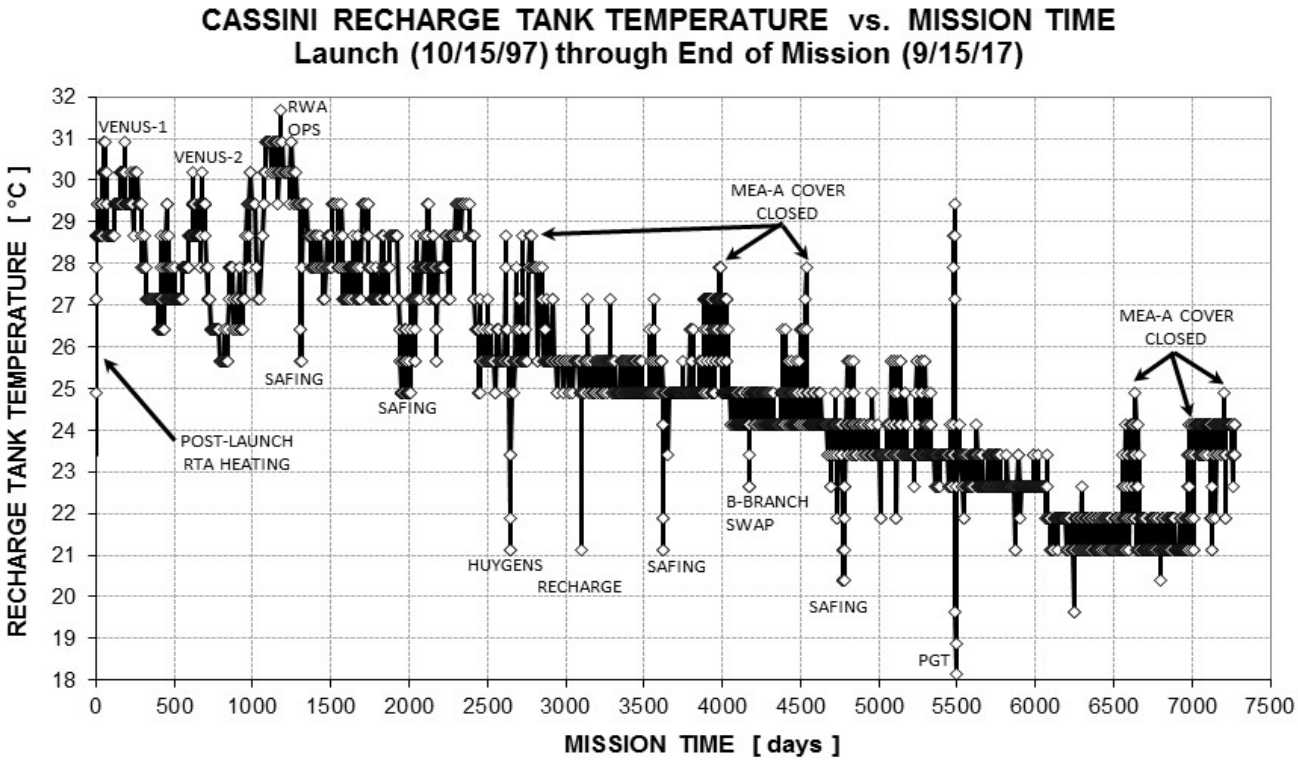

Figure 53: Cassini Recharge Tank Temperature vs. Mission Time

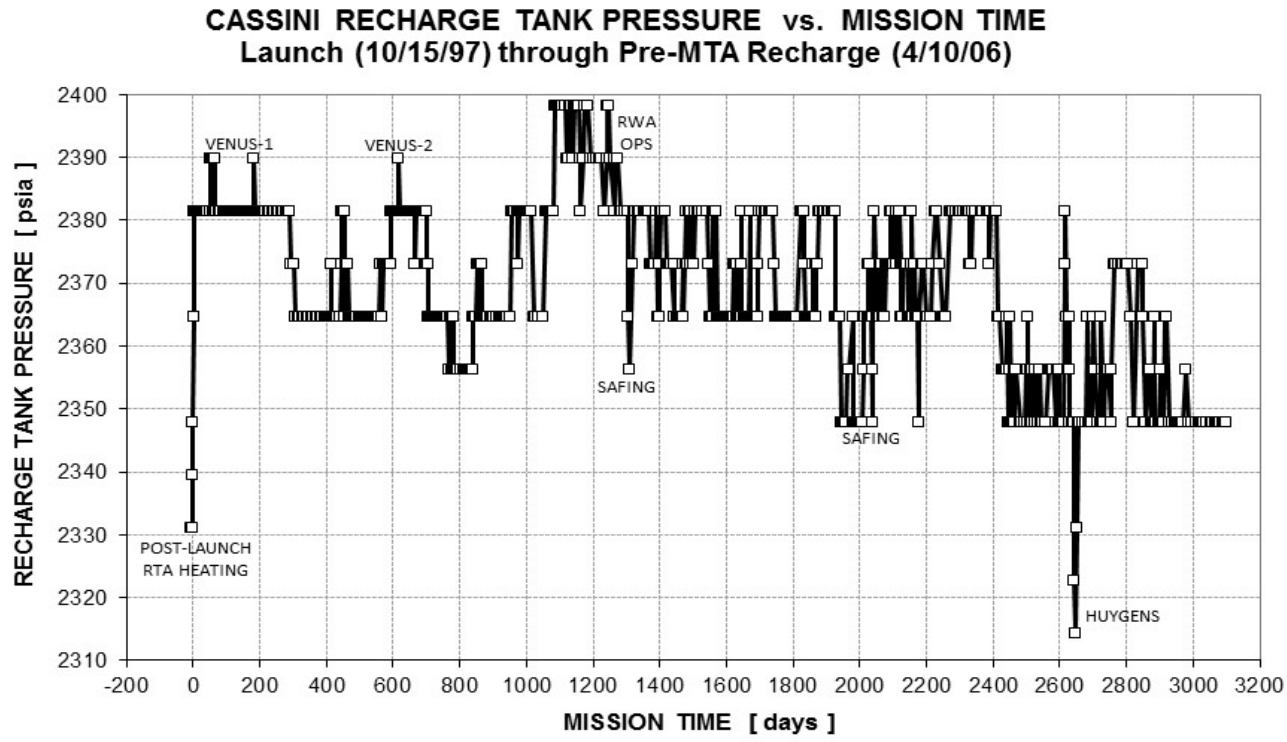

Figure 54: Cassini Recharge Tank Pressure vs. Mission Time (Launch thru Pre-MTA Recharge)

In theory, after recharge the RTA pressure and hydrazine tank pressure telemetry should agree exactly, since PV40 firing connected the tanks and thus equilibrated pressures for the remainder of the mission. However, one key distinction is the DN resolution of these sensors. Since RTA pressure was as high as 2400 psia before recharge, a high-range transducer was required, leading to rather crude DN resolution of 16.8 psia/DN. In contrast, the N_2H_4 tank transducer resolution was much finer, about 1.88 psia/DN. Therefore, agreement no better than about 20 psid should be expected between these two transducers, even when measuring the same physical pressure. In fact, following recharge, these two sensors maintained a difference of ~20 psid for the rest of the mission, with the RTA reporting lower pressure than the hydrazine tank.

If the lower values measured by the RTA pressure were to be believed, this may have been a cause for concern during the mission, since this would have implied there was quite a bit less remaining N_2H_4 than modeled. However, it is much more likely the hydrazine tank transducer measured true N_2H_4 pressure, given its much finer DN resolution. Figure 55 presents these two pressure telemetry curves as a function of mission time following recharge, through the end of mission. Other than the explainable difference between the two curves, the only other notable features in Figure 55 are minor pressure transients due to spacecraft spacing, a more substantial change caused by PGT, and a Solstice mission which was rather miserly with respect to hydrazine usage (fortunately).

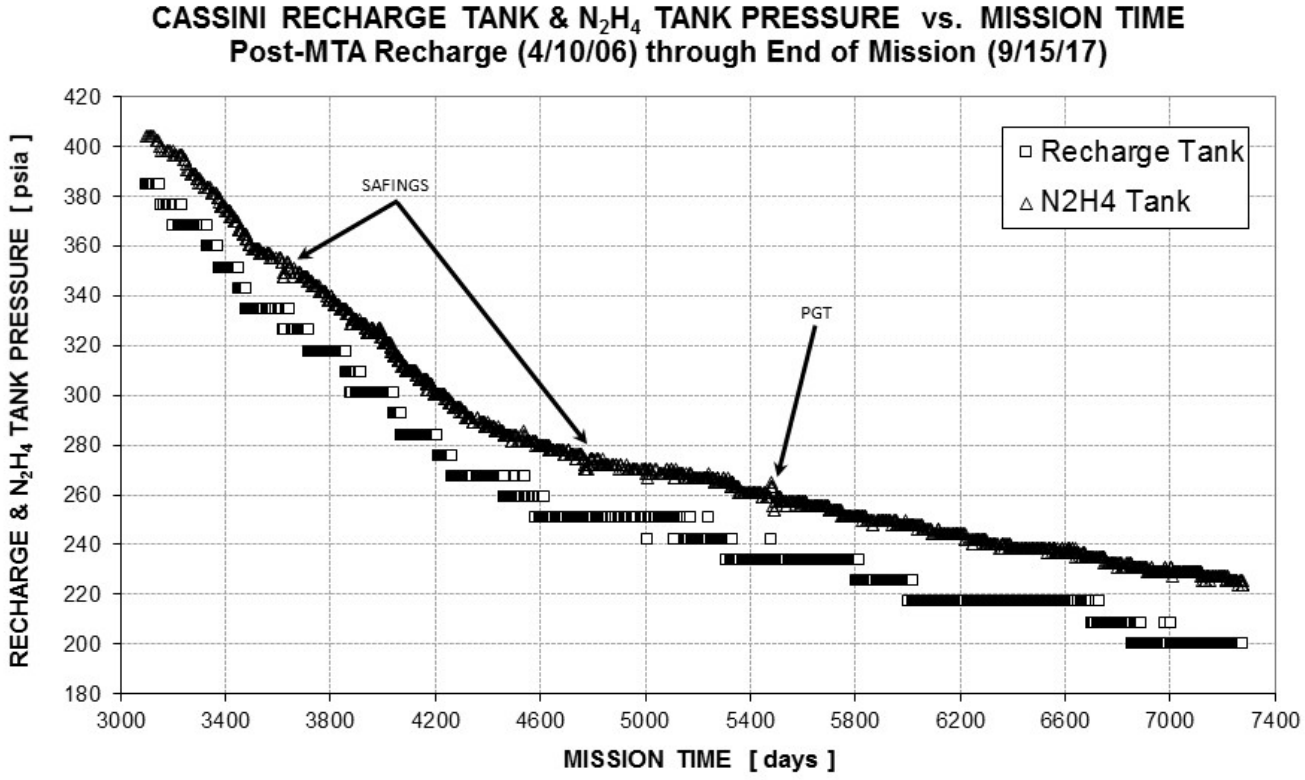

Figure 55: Cassini Recharge Tank and N_2H_4 Tank Pressure vs. Mission Time (Recharge thru End of Mission)

Focusing now just on hydrazine tank temperatures and pressure telemetry for the entire mission, Figure 56 presents the liquid hydrazine and ullage gas temperatures from launch through Saturn plunge. There is very nearly a one-to-one correspondence between Figure 56 and Figure 53, with the labeled spacecraft activities causing quite similar transients in the N_2H_4 tank as they did in the RTA. From the CPMS layout presented in Figure 7, these two tanks were in close proximity to each another, so this agreement is perhaps not surprising.

One item of note from Figure 56 is a persistent thermal gradient within the hydrazine tank, over the entire mission. Liquid hydrazine temperatures were typically 3-4°C warmer than gas-side temperatures while the main-engine cover was open, decreasing to roughly 2-3°C while the cover was closed. This minor thermal gradient was

largely due to RTG waste heat warming the N_2H_4 tank from the bottom. Even though the difference between gas and liquid side temperatures was well within requirements, a post-launch concern was identified given a potential "cold trap" at the tank gas inlet. Specifically, it was possible N_2H_4 vapor could be permanently lost across the AF-E-332 diaphragm due to permeation, diffusion, and condensation at the cold trap. Unfortunately, there was no way to reverse the thermal gradient in Cassini's hydrazine tank, so whatever hydrazine mass was lost due to this "pumping" mechanism was rendered unusable monopropellant. In fact, this non-negligible contributing factor for yet more unusable hydrazine will be covered in the next section. In-flight ranges for N_2H_4 tank liquid hydrazine and ullage temperatures ended up being 15-29°C and 13-26°C, respectively, within the AFT (operating and non-operating) limits of 10-45°C, comfortably so on the upper end.

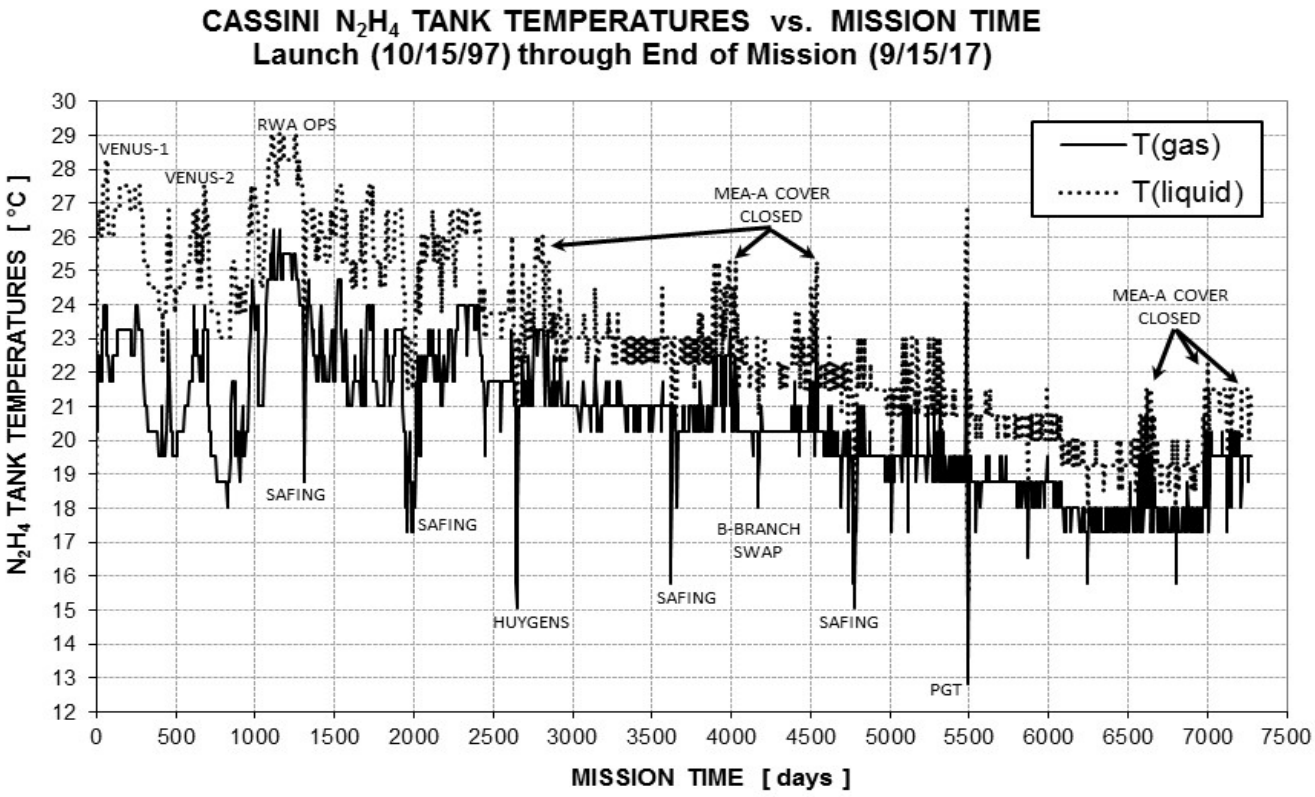

Figure 56: Cassini N_2H_4 Tank Hydrazine Liquid & Ullage Gas Temperatures vs. Mission Time

Figure 57 presents N_2H_4 tank pressure over the entire mission, now using data from just the low-range (fine) transducer. This is a classic blowdown pressure curve, interrupted only by the successful recharge midway through the prime orbital tour at Saturn. Without recharge, control authority would have been insufficient to fight Titan's atmosphere during subsequent low Titan flybys, so its success was paramount. The only other items of note in Figure 57 are minor tank pressure transients due to N_2H_4 tank temperature changes around spacecraft safings, PGT, the Huygens mission, and RWA operations 1999-2000. The as-flown range for hydrazine tank pressure was 223-404 psia, securely within the MR-103H thruster inlet pressure allowable firing range of 90-420 psia.

Assessing now the three tanks within the bipropellant portion of the CPMS, Figure 58 below displays the HTA temperature over the entire Cassini mission. Even though the HTA was on the opposite side of the orbiter vs. the RTA and MTA (see Figure 7), it is remarkable how similar Figure 58 is to Figure 53 (and, to a lesser extent, Figure 56). The post-launch thermal transient discussed previously is visible in Figure 58, the one that caused a permanent jog in the bipropellant system helium budget. The only transients in Figure 58 which were not evident in Figure 53 are relatively large thermal transients associated with DSM, SOI, and PRM polytropic cooling in the HTA and the

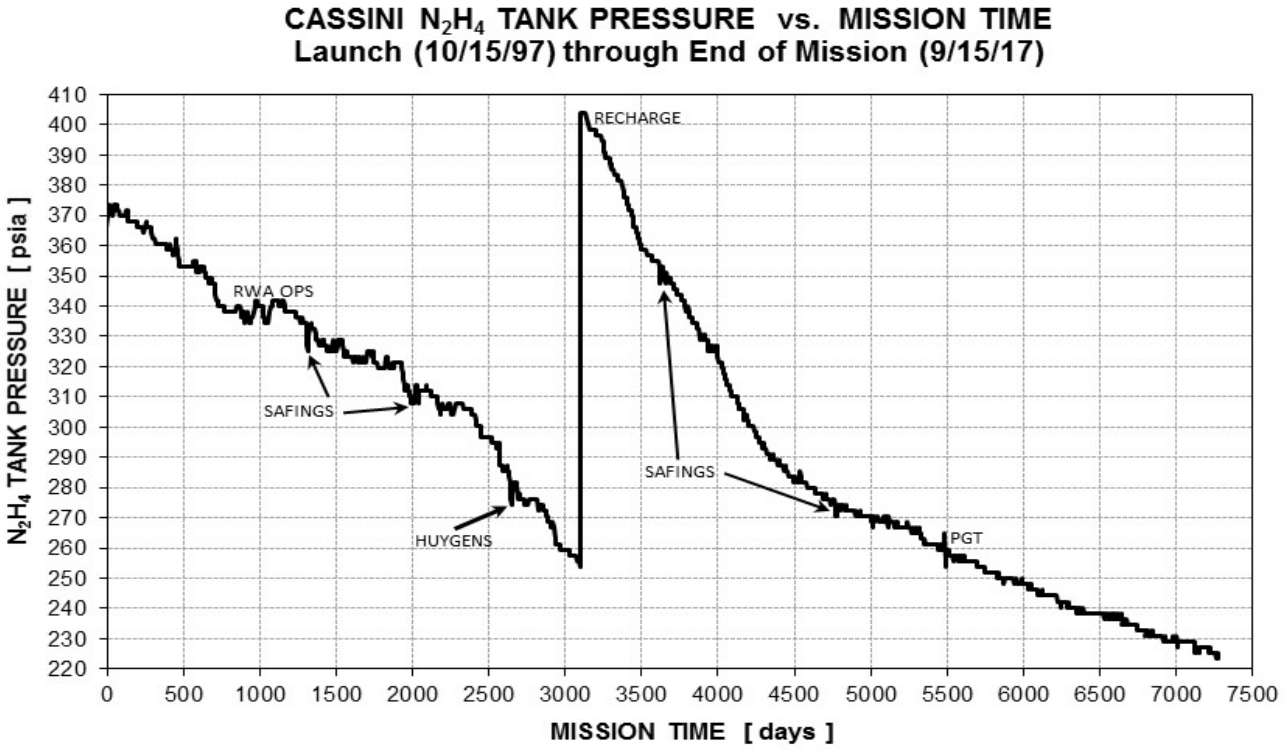

Figure 57: Cassini N_2H_4 Tank Pressure vs. Mission Time

constant decrease in HTA temperature caused by turning off PCA panel and helium line heaters in 2014. Though surely the actual minimum HTA temperatures during regulated burns were likely not captured in telemetry, the in-flight range over nearly two decades was 10-30°C, just within the AFT limits of 10-45°C at the lower end. The minimum value of 10°C occurred during DSM and is off-scale in Figure 58, in order to show finer detail over the entire mission. Despite these relatively tight AFT limits, transient HTA temperatures as low as -40°C were permissible in flight, so there was never any concern for helium tank temperatures during Cassini's flight.

HTA tank pressure telemetry over the mission is presented below in Figure 59. Changes in HTA pressure were dominated by the nine helium transfer events mentioned previously, rather than by the relatively minor changes in HTA temperature seen in Figure 58. As such, Figure 59 strongly resembles the inferred helium mass history from Figure 34, as expected. The only thermal transients which caused even as much as a wiggle in Figure 59 were from polytropic cooling during regulated burns (particularly DSM) and tank heating and cooling during PGT. HTA pressure in flight spanned a range of 870-3458 psia, and both Figure 34 and Figure 59 were used to confirm the pressure transducer in the HTA remained drift-free over this lengthy mission.

As with the N_2H_4 tank, it was possible to measure the gas and liquid temperatures separately in the NTO and MMH tanks. In fact, there were six temperature sensors provided for each bipropellant tank, TO1-TO6 and TF1-TF6, spatially distributed around the tanks (see Figure 8). Except near the end of mission (i.e., with nearly empty NTO and MMH tanks), a reasonable estimate of gas side temperature was the average of TO1 & TO4 for the NTO tank and TF1 & TF4 for the MMH tank. Similarly, bipropellant liquid average temperatures were calculated as the average of the four remaining sensors, (TO2+TO3+TO5+TO6)/4 or (TF2+TF3+TF5+TF6)/4. For clarity, a separate plot will be provided below for gas and liquid average temperatures by tank.

Figure 60 displays the NTO tank average gas temperature over the mission. Even though both the NTO tank and MMH tank were embedded within the central body of the orbiter, it is again remarkable how well temperature followed other CPMS tanks on the spacecraft. All of the previously described spacecraft activities which caused thermal transients affected the NTO gas temperature similarly, as may be seen in Figure 60. Some activities actually elicited a more pronounced response than witnessed in other tanks, including the gas cooling during regulated burns (most notably DSM), main-engine cover cycling, and PGT.

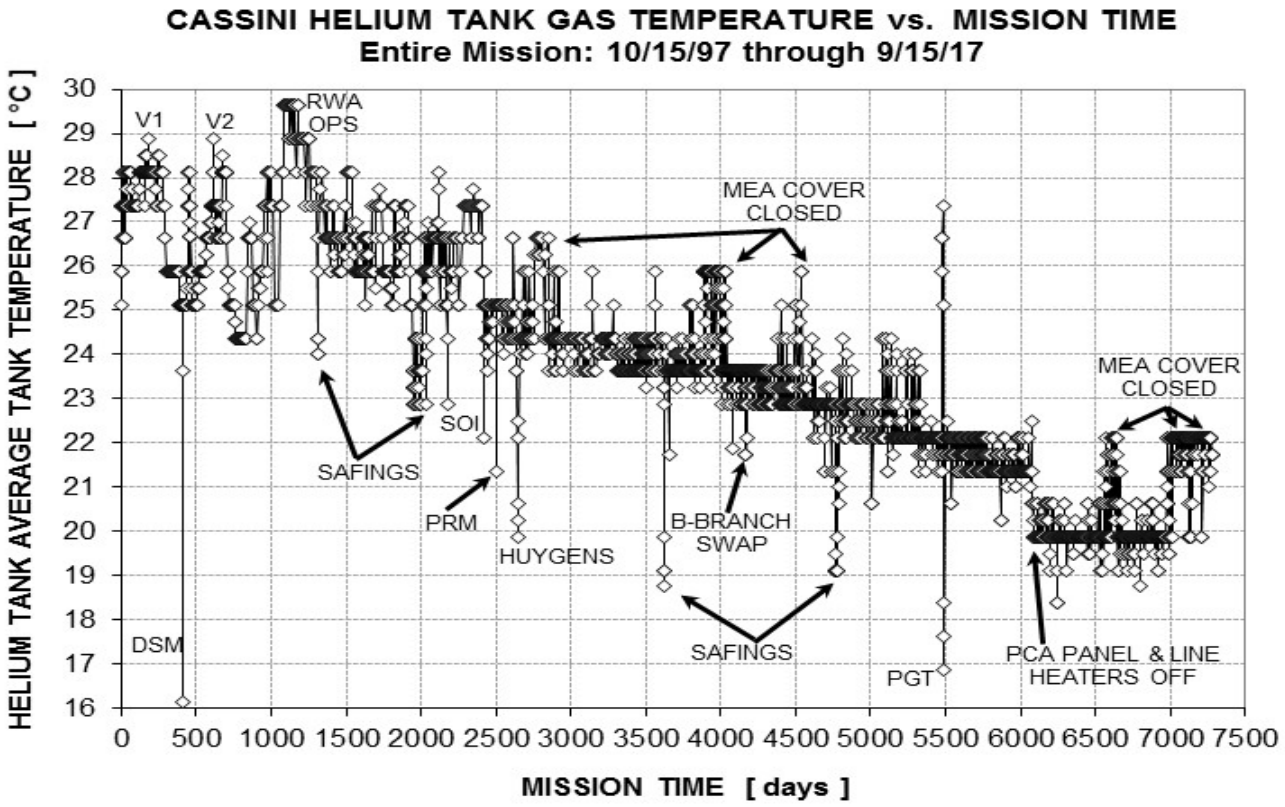

Figure 58: Cassini Helium Tank Temperature vs. Mission Time

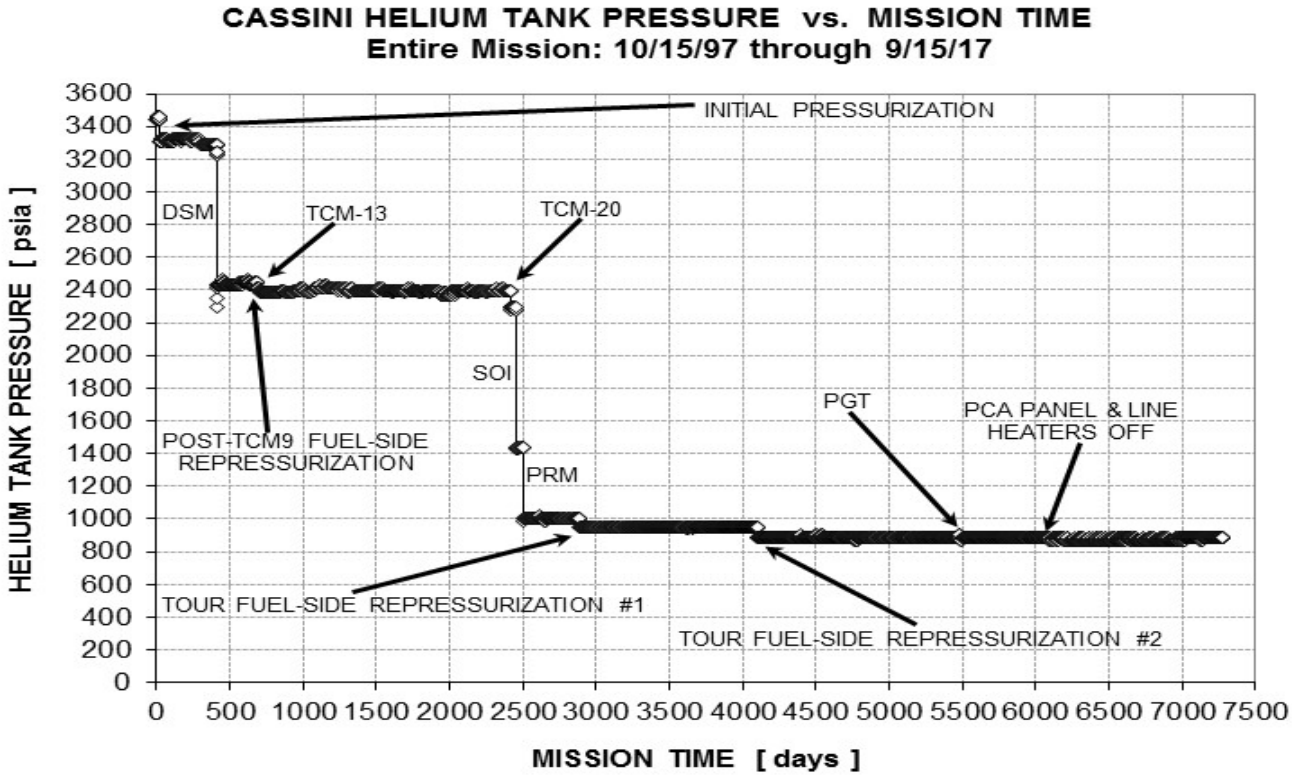

Figure 59: Cassini Helium Tank Pressure vs. Mission Time

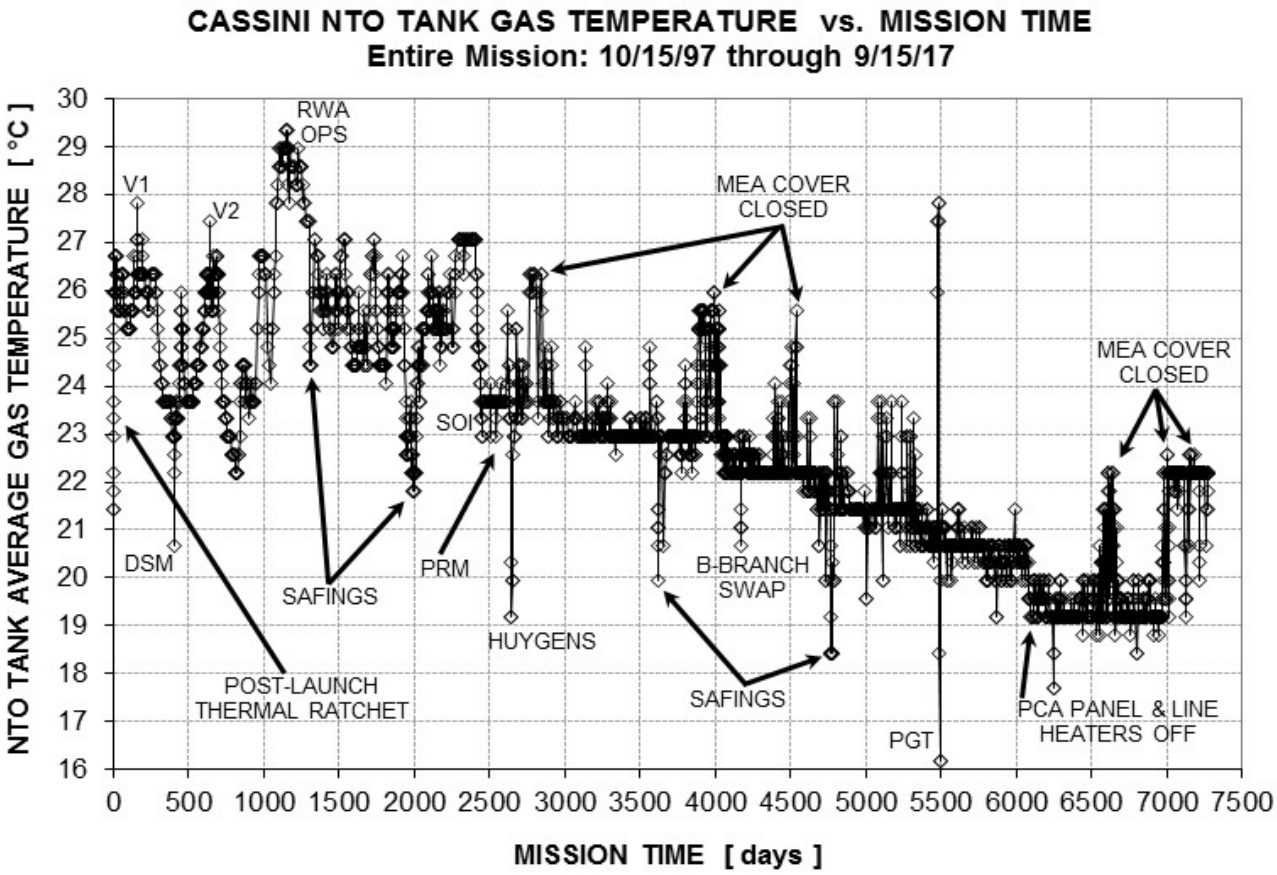

Figure 60: Cassini NTO Tank Average Gas Temperature vs. Mission Time

One spacecraft activity not previously discussed is plainly visible in Figure 60, a sharp temperature increase just after launch. This "thermal ratcheting" event was a planned turn-on of the NTO and MMH tank heaters, executed in order to perform initial pressurization at higher temperatures, thus preventing tank overpressurization near perihelia and the Venus flybys. Unfortunately, bipropellant tank heaters each emanated a rather sporty 37 W of heat directly into the liquid sides of the NTO and MMH tanks, causing liquid temperatures to rise far faster than gas temperatures. In fact, within a few days, the allowable gas/liquid thermal gradient across the NTO tank (8°C) was exceeded, so real-time commands were sent to Cassini to turn off bipropellant tank heaters. Thankfully, enough thermal ratcheting persisted, even after tank heaters were turned off, to preclude any overpressurization problems during VVEJGA. Incidentally, tank heaters were never turned on again, though their use during PGT was contemplated. Thermal gradient issues with nearly empty tanks would occur much more quickly, of course, so this approach had to be abandoned. This is unfortunate, because the direct heating available from tank heaters would have been ideal for PGT.

Comparing in-flight NTO gas temperature range with AFT, margins were adequate vs. requirements. As may be seen in Figure 60, the oxidizer gas remained within a range of 16-29°C during the twenty-year Cassini mission. This was comfortably within the non-operating AFT limits of 5-45°C and symmetrically nestled within the operating (i.e., at-ignition or firing) AFT limits of 10-35°C. These same AFT limits apply for NTO liquid temperature, so the non-operating limit of 35°C was actually exceeded slightly (however briefly) during thermal ratcheting, as may be seen below in Figure 61, the time history of NTO liquid temperatures. Fortunately, conditions were well within acceptable limits by the time of TCM-1, a few days after bipropellant tank heater turn-off. The in-flight range of values for NTO liquid temperature were 16-36°C, and NTO liquid and gas temperature time histories were very similar (cf. Figure 61 and Figure 60). The most notable difference between Figure 61 and Figure 60 is the muted thermal transient response in NTO liquid temperature (as expected, given thermal inertia) vs. gas-side temperatures. A more subtle difference is the increased telemetry resolution in Figure 61 vs. Figure 60; this is also as expected, since four sensors were averaged to obtain the data of Figure 61 vs. only two in Figure 60.

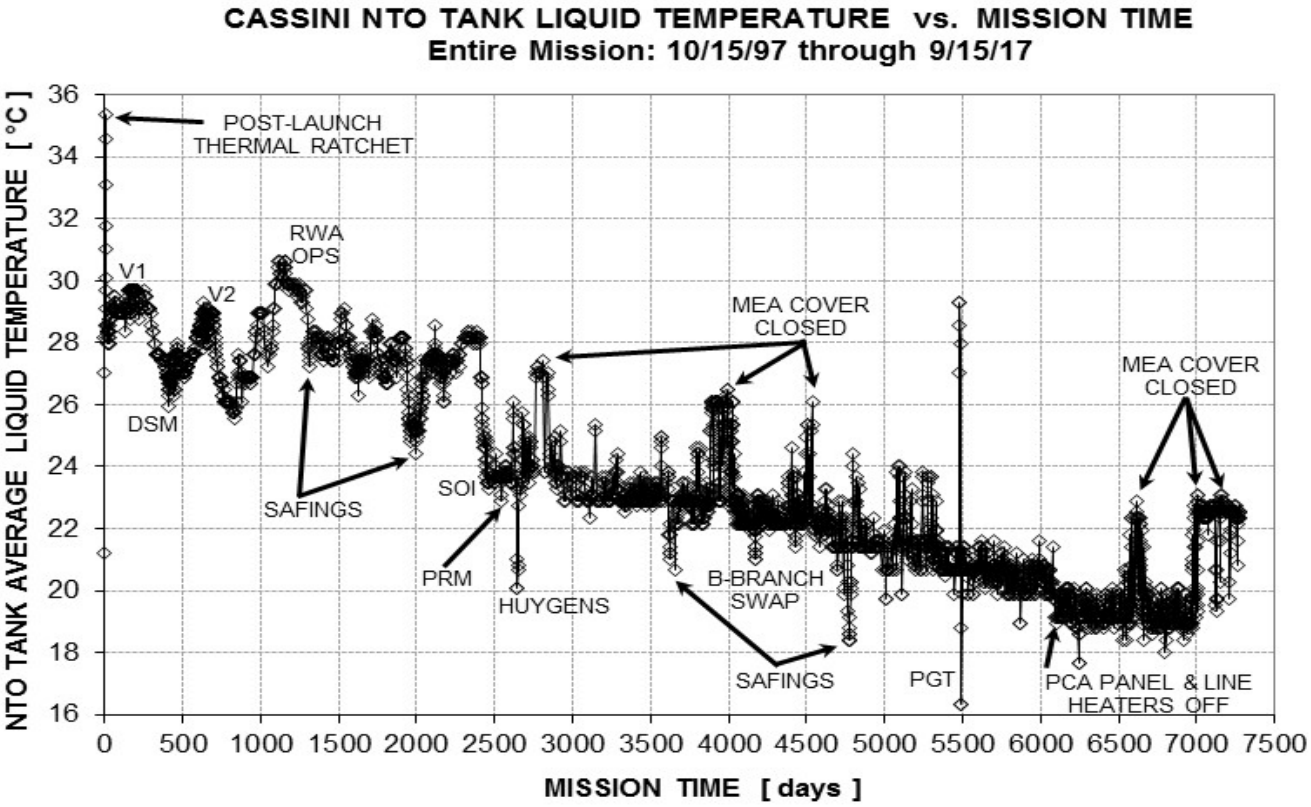

Figure 61: Cassini NTO Tank Average Liquid Temperature vs. Mission Time

The MMH analogue to Figure 60 is presented below as Figure 62. Note how similar MMH and NTO gas temperatures were throughout nearly two decades of spaceflight. Minor differences between these two quantities included slightly warmer conditions for the NTO tank vs. the MMH tank (as expected, given RTG locations) and larger transients in the MMH tank during times of low spacecraft power usage (safings, B-branch swap, and Huygens). This is also as expected, since the MMH tank was closer to and thus more closely coupled to central body temperature than the NTO tank. The in-flight range of MMH gas temperature was 15-27°C, again well within the non-operating AFT limits of 5-45°C and never approaching the operating (i.e., at-ignition or firing) AFT limits of 10-35°C.

Figure 63 presents the MMH tank liquid temperature history over the mission. As expected, it closely resembles Figure 61, the equivalent plot for the NTO tank. Post-launch thermal ratcheting did not cause as large a temperature gradient in the MMH tank between the gas and liquid sides, only 5°C vs. 8°C in the NTO tank. Still, the MMH tank heater was turned off at the same time as the NTO tank heater, since it was only a matter of time before its gradient limit of 8°C were also breached. The in-flight span of MMH liquid temperatures was 15-32°C, vs. the non-operating AFT limits of 5-45°C and operating (i.e., at-ignition or firing) AFT limits of 10-35°C. Controlling CPMS temperatures for two decades was no small feat, given wide ranges in solar distance and copious spacecraft activities which affected power margin. Kudos are due Cassini thermal engineers for enabling such benign conditions for the CPMS.

To conclude this section, Figures 64 and 65 below capture the NTO and MMH tank pressure history during the entire Cassini mission. At launch, bipropellant tank pressures were rather low, very close to 100 psia. Post-launch heating brought NTO and MMH tank pressures up to 114 and 103 psia, respectively, just before initial pressurization. To discern finer detail in Figures 64-65, the vertical scale of these plots was selected to span the range of pressures encountered after initial pressurization (i.e., initial pressures are off-scale low in the figures below).

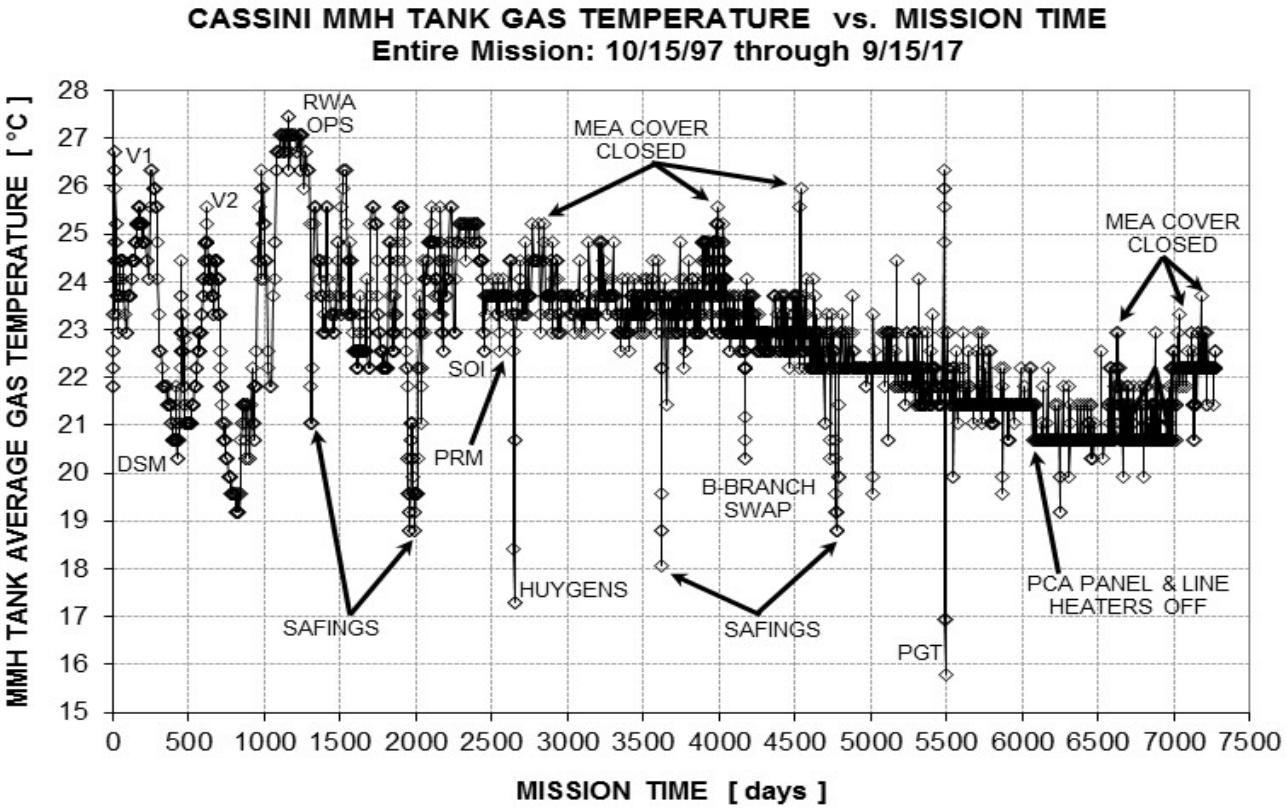

Figure 62: Cassini MMH Tank Average Gas Temperature vs. Mission Time

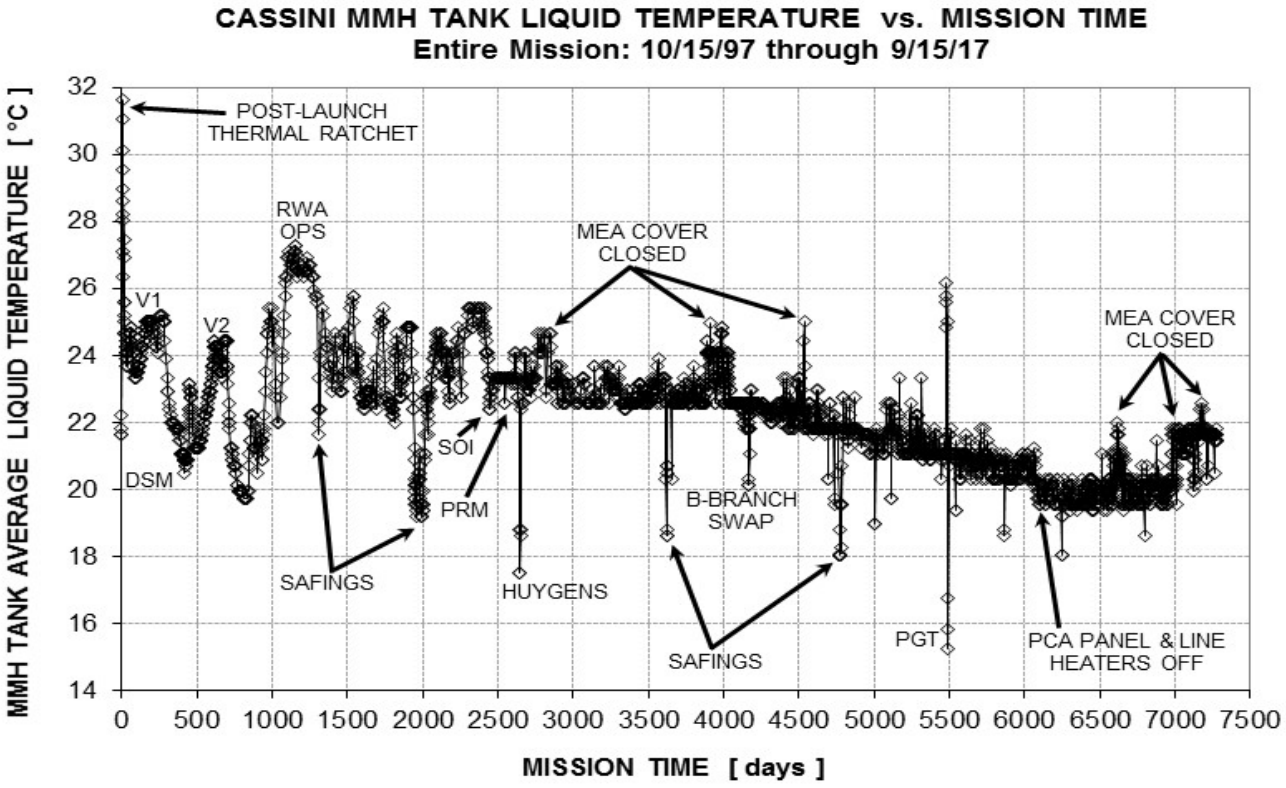

Figure 63: Cassini MMH Tank Average Liquid Temperature vs. Mission Time

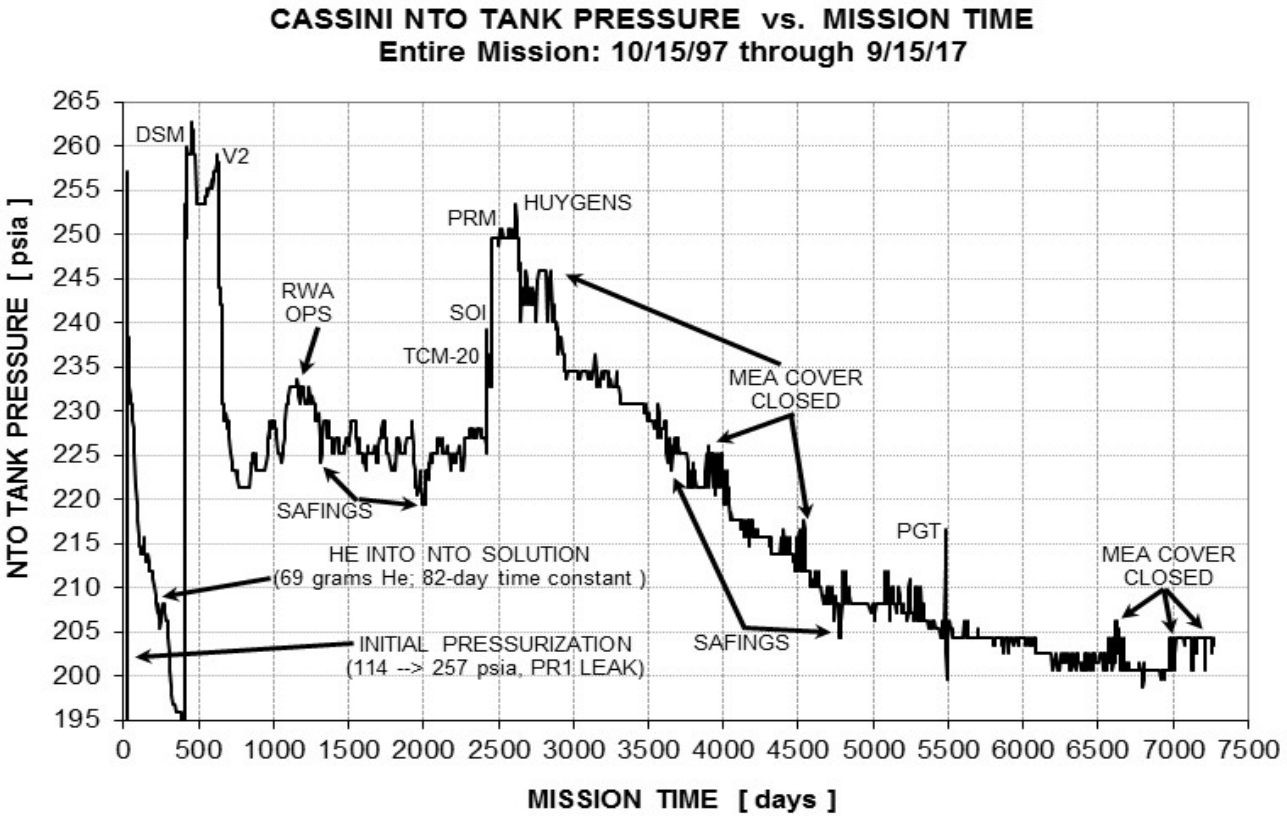

Figure 64: Cassini NTO Tank Average Pressure vs. Mission Time

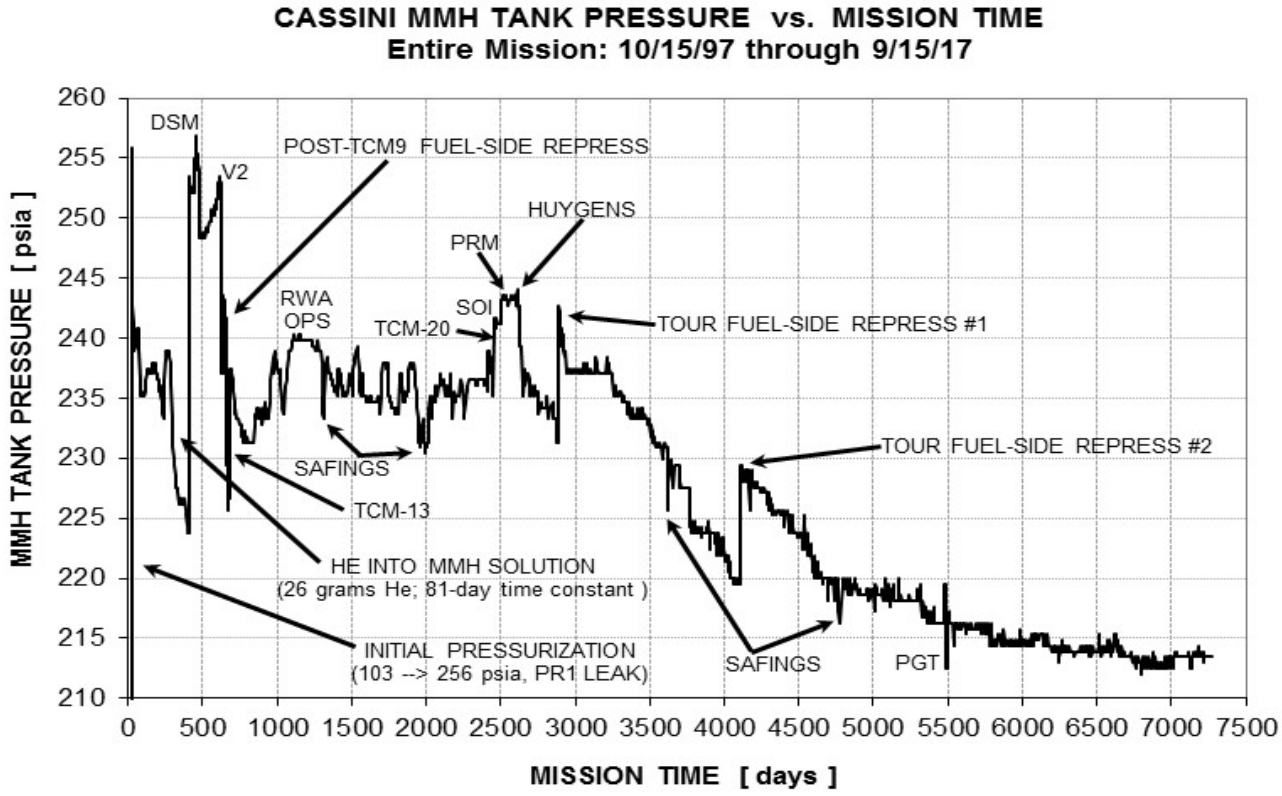

Figure 65: Cassini MMH Tank Average Pressure vs. Mission Time

The pressure traces in Figures 64-65 are rather complex, since NTO and MMH tank pressures were influenced by (1) helium solubility after initial pressurization, (2) tank temperature changes, and (3) bipropellant pressurization events. Between these events, however, the curves follow classic blowdown propellant tank decay curves, as expected. Specific tank pressure conditions during main-engine firings were presented in Figure 9, the operating box (NTO tank pressure vs. MMH tank pressure) flight history for Cassini's 183 R-4D firings. Hopefully this complete set of CPMS tank pressure and temperature plots covering the entire 1997-2017 time frame will aid future projects as inevitable questions arise during design and development of these prospective missions.

XII. Unusable Propellant Assessment

By now it is abundantly clear one of the biggest concerns during the mission, particularly the extended missions, was propellant depletion. Bipropellant margins, especially, were razor thin during the last few years of the Solstice mission. Typically, pre-launch values for unusable propellant are conservative, since they must encompass a wide variety of off-nominal conditions which may never materialize in flight. Often, unusable propellant numbers are scrubbed post-launch, using as-flown conditions for propellant consuming events. This invariably reduces unusable propellant, allowing further resources to be available for extended missions. Unusually, Cassini had the opposite experience—it seemed the more the propulsion team poked at unusable propellant, the larger the values grew!

Across a number of missions, there seems to be consistent values specified for unusable propellant. For hydrazine and NTO/MMH bipropellant systems, typical 3σ unusable percentages (deterministically plus statistically unusable propellant) have been 1.0% and 2.7%, respectively. Galileo baselined the latter value for its bipropellant system, but it was able to be reduced during flight as actual firing conditions became known. Given Cassini's large launch loads of 1869 kg of NTO and 1131 kg of MMH, 2.7% unusable bipropellant translates to 50.5 kg of NTO and 30.5 kg of MMH. Contrast these numbers with the final reckoning of NTO and MMH masses at plunge, including the -2.7 kg NTO correction and -5.9 kg MMH correction from Section IX. With these adjustments, final Cassini bipropellant numbers at end of mission were 35.6 kg of NTO and only 14.9 kg of MMH. Note these are well below the pre-launch unusable bipropellant uncertainties, but what matters is post-launch uncertainty.

Table 20 summarizes the post-launch scrub of unusable bipropellant for Cassini. This work began in earnest in 2008, though an additional error term was only recognized in 2013, during the independent propulsion assessment of PGT and bipropellant uncertainties. To the point, there was an unaccounted ±2% analog-to-digital (A-to-D) uncertainty for all analog channels. Presumably this was a 3σ spec value, but its effect on tank pressure error was devastating. Based on experience, Cassini engineers deemed a ±2% A-to-D error way too pessimistic, so the Cassini Spacecraft Operations Team Chief recommended a 3σ spec value of ±1% A-to-D error be used instead. Even with this reduced error, the tank pressure uncertainty line in Table 20, for both NTO and MMH, doubled vs. the initial values from 2008. Clearly, this inherent error was a dominant factor in calculating unusable bipropellant.

Taking each row in Table 20 by turn, the first three rows are contributors to deterministically unusable bipropellant. A more traditional name for this source of unusable propellant might be hold-up, propellant physically unable to be consumed. From line volume measurements downstream of the propellant tanks, along with knowledge of line temperature and propellant densities, the trapped liquid NTO and MMH masses in the lines were determined to be only 0.7 and 0.4 kg, respectively. Another contributor to hold-up propellant is the vapor mass in the NTO and MMH tanks. For the fuel tank, it was a mere 0.2 kg, given MMH's low vapor pressure. In contrast, an impressive 5.9 kg of NTO vapor would remain trapped in the NTO tank at liquid oxidizer depletion. The hold-up term was due to the BTA expulsion efficiency of 99.5%. Adding these trapped propellant masses led to the total hold-up or deterministically unusable bipropellant, 15.9 kg of NTO and 6.3 kg of MMH. These numbers are already ~40% of the final total remaining NTO and MMH masses at plunge, not an encouraging start.

In theory, there are two ways to calculate statistically unusable bipropellant, either by (1) determining errors in consumed NTO and MMH over 183 main-engine burns or (2) calculating errors in the final NTO and MMH tank conditions based on tank model uncertainties. Unfortunately, the large error bars associated with nine helium transfer events from the HTA to the BTA render the second method quite inaccurate, so modeling consumption errors was the best way to minimize unusable NTO and MMH. Starting with a quite simple term, there were minor propellant loading uncertainties at Cape Canaveral, and these were known before launch, ±3.4 kg of NTO and ±2.2 kg of MMH (3σ).

Table 20: Cassini Unusable Bipropellant Mass Summary Table

Deterministically Unusable Bipropellant	**NTO [kg]**	**MMH [kg]**	**NTO%**	**MMH%**
Trapped Liquid in Propellant Lines	0.7	0.4	0.04	0.04
Trapped Vapor in Propellant Tank @ Depletion	5.9	0.2	0.32	0.02
Trapped Due to 99.5% Tank Expulsion Efficiency	9.3	5.7	0.50	0.50
Total Deterministic	15.9	6.3	0.85	0.56
Statistically Unusable Bipropellant (3σ)	**NTO [kg]**	**MMH [kg]**	**NTO%**	**MMH%**
Bipropellant Loading Uncertainty	3.4	2.2	0.18	0.19
Tank Pressure Uncertainty	27.0	17.2	1.44	1.52
Line Pressure Drop Uncertainty	7.7	3.3	0.41	0.29
Tank Temperature Uncertainty	1.6	0.4	0.09	0.04
Main-Engine Burn Time Uncertainty	0.6	0.3	0.03	0.03
Delivered Mixture Ratio Uncertainty	6.6	4.0	0.35	0.35
Total Statistical (3σ, RSS)	29.1	18.1	1.56	1.60
TOTAL UNUSABLE PROPELLANT	45.0	24.4	2.41	2.16

The next row in Table 20 was not only the largest contributor to statistically unusable bipropellant, it was also the most complex, being comprised of many individual sources of error. Tank pressure uncertainty included 3σ spec uncertainties for linearity (±0.5%), hysteresis (±0.2%), and repeatability (±0.1%). However, the largest contributor by far was the compensated temperature range spec of ±2.0% (3σ), a true killer for the unusable bipropellant budget. Fortunately, though, this spec was conservative because it had to apply over the entire qualification temperature range of -40°C to +45°C. However, since flight BTA temperatures (other than during brief transients) varied within a very tight range of only 19-27°C, the compensated temperature range spec of ±2.0% (3σ) was linearly scaled by the actual temperature range in flight divided by the qual range. This reduced the ±2.0% (3σ) spec to only ±0.19% (3σ). These terms were then RSS'd with the assumed worst case ±1% A-to-D error, along with the ± ½ DN uncertainty of the tank pressure itself. Finally, this RSS'd pressure measurement error was converted to NTO and MMH mass error at bipropellant depletion, leading to the largest terms in Table 20, 27.0 kg of NTO and 17.2 kg of MMH. Incidentally, the spec for pressure transducer drift is ±0.1% per year, up to ±2.0% total (presumably 3σ). Given truly miniscule Cassini BTA pressure transducer drift during flight, this separate term was not included in Table 20. It was assumed to be covered within existing pressure transducer error budgets.

Line pressure errors during main-engine firing also contributed to unusable bipropellant; these 3σ errors are typically taken as 5% of the line pressure drop or 1 psid, whichever is larger. For the Cassini REA-A lines, NTO and MMH pressure drops were about 30 psid and 20 psid, respectively, suggesting line pressure error should be taken as ±1.5 psia for NTO and ±1.0 psia for MMH. This translated into 7.7 kg and 3.3 kg of 3σ statistically unusable NTO and MMH, respectively, as may be seen in Table 20. Tank temperature uncertainties were likely conservative, using the sum (not RSS) of DN uncertainties and a known offset present in Solar Thermal Vacuum (STV) testing. Finally, for each of 183 main-engine burns on REA-A, minor errors in actual burn time and more substantial potential errors in as-flown mixture ratio introduced two more rows in Table 20. The six terms contributing to statistically unusable bipropellant were then RSS'd to determine 3σ totals for statistically unusable NTO and MMH, 29.1 kg and 18.1 kg, respectively. Note how these numbers are just barely above the unusable numbers due to tank pressure uncertainties themselves, a demonstration of the mitigating effect of RSSing.

As a last step, the deterministically and statistically unusable bipropellant terms were added to form the total 3σ unusable NTO and MMH, the final row in Table 20. 45.0 kg of NTO (2.41% of the launch load) and 24.4 kg of MMH (2.16% of the launch load) were considered unusable; these values were also demarcated on Figure 32 above. One piece of good news is the unusable bipropellant percentages, 2.41% for NTO and 2.16% for MMH, were less than the pre-launch placeholder value of 2.7%. This did little to assuage fears of bipropellant depletion, however. In fact, it is now possible to retroactively calculate the probability that Cassini "should" have statistically run out of

NTO or MMH before plunging into Saturn. With a final NTO mass of 35.6 kg, 15.9 kg of that hold-up, there was only about 19.7 kg of usable NTO remaining in the mean. However, statistically unusable NTO was ±29.1 kg (3σ), so a 2.03σ "bad day" would have depleted NTO before end of mission. From the normal distribution table, this would be expected to occur about 2.1% of the time. A similar calculation may be performed for unusable MMH. With a final MMH mass of only 14.9 kg, 6.3 kg of that deterministically unusable, there was only about 8.6 kg of usable MMH remaining. However, statistically unusable MMH was ±18.1 kg (3σ), so a 1.43σ case or worse would have depleted MMH before end of mission. Again from the normal distribution table, this would be expected to happen about 7.6% of the time. Therefore, assuming these errors are independent (only mixture ratio uncertainties couple these two errors), the probability of not running out of NTO or MMH was (1-0.021)·(1-0.076) = 90.5%, or, equivalently, the probability of depleting NTO or MMH was about 9.5%! The author is grateful this calculation was first performed eight months after the plunge, explicitly for this paper, given the sobering results.

Focusing now on unusable hydrazine, as before, there were two ways to calculate statistically unusable N_2H_4, either by (1) determining errors in consumed hydrazine over the mission or (2) calculating errors for an empty monopropellant tank based on tank model uncertainties. Unlike the bipropellant case, lower error bars were possible by using (2) rather than (1), largely because hydrazine consumption errors were quite large. Table 21 below is the N_2H_4 analogue to Table 20, and some of the rows are equivalent for both tables. Looking first at deterministically unusable N_2H_4, the trapped liquid in propellant lines (0.608 kg) and trapped hydrazine vapor mass at N_2H_4 depletion (only 0.004 kg) are precisely the same terms as the first two rows of Table 20. The fourth row is equivalent as well, the term due to 99.5% tank expulsion efficiency. Notably, though, the N_2H_4 tank expulsion spec was 99.0% (3σ). Through conversations with JPL/LMA propulsion experts and Pressure Systems, Inc. (PSI) engineers, including personnel who actually built Cassini's hydrazine tank, expulsion efficiencies of 99.7-99.9% or better were more typical. This allowed the Cassini propulsion team to relax this requirement.

There were three additional terms in Table 21 for hold-up N_2H_4, ones not seen in Table 20. A couple of deterministic errors were due to tank temperature sensor offset (as noticed in STV testing) and differences between line and tank pressure (see Figure 20). However, the most interesting deterministically unusable term in Table 21 is in the third row. During a Mars rover design review, the author noted a new JPL requirement to keep gas-side temperatures warmer than liquid-side temperatures in N_2H_4 tanks with diaphragms. In theory (and based on anecdotal evidence), with cooler gas-side temperatures, there was a mechanism for permanent N_2H_4 loss. Specifically, N_2H_4 vapor permeation through the diaphragm, followed by diffusion and local condensation at a cold trap, could lead to loss of usable liquid N_2H_4. If Mars missions lasting less than a year were concerned about this, Cassini was potentially in trouble, given its two-decade mission, persistent unfavorable thermal gradient (see Figure 56), and inability to reverse this thermal gradient. The author developed a steady-state model, based on known temperature differences, measured N_2H_4 permeation rates through AF-E-332 rubber, and Fick's law. Fortunately, even over the long Cassini mission, no more than about 1 kg of hydrazine could have been lost via this mechanism.

Summing the six contributors to deterministically unusable N_2H_4 from Table 21, hold-up monopropellant was about 3.65 kg, or 2.76% of the launch load of 132.0 kg. This is already quite a bit higher than the pre-launch assumption of 1.0% unusable N_2H_4. Worse, statistically unusable hydrazine has not even been considered yet. It should be noted, however, the three additional error terms covered in the prior paragraph were new to Cassini and thus not covered in any pre-launch assessment of hold-up N_2H_4. Omitting these terms, hold-up hydrazine would have only been 1.272 kg, 0.96% of the launch load and thus quite close to the traditional value of 1.0%.

The terms which make up statistically unusable N_2H_4 included propellant loading uncertainty, tank pressure uncertainty at depletion (both from standard errors and from transducer drift), and uncertainties in tank temperature and helium mass at the hydrazine depletion. Loading uncertainty was calculated to be ±0.911 kg (3σ), based on conversations with Lockheed Martin Astronautics engineers at Cape Canaveral for propellant loading. Skipping tank pressure errors for the moment, uncertainties in tank temperature and total helium mass in the RTA and MTA combined tank volume also contributed to hydrazine mass errors in the tank model. Fortunately, errors in tank temperature had only second-order effects on calculated N_2H_4 mass vs. the first-order effects from pressure error. Moreover, the fixed offset from STV data was already included in deterministically unusable N_2H_4, so the only uncertainty which needed to be included was the ± ½ DN uncertainty of the tank temperature sensors themselves. Indeed, this turned out to be a small contributor to statistically unusable N_2H_4, only ±0.531 kg (3σ).

Determining the error in monopropellant system helium mass was rather simple, since final helium mass was the sum of the loaded RTA & MTA helium masses minus the tiny amount of helium in N_2H_4 solution expelled to space. Uncertainties for these three contributors were calculated using Figures 30-31 and an assumed 3σ helium solubility uncertainty of ±30%, respectively. The latter is a generous error bar, but helium solubility results in NTO and MMH after initial pressurization suggest applying some conservatism here. Running the Cassini tank model for a depleted tank revealed the contribution from final helium mass uncertainty, and it turned to be only ±0.869 kg (3σ).

Tackling now tank pressure uncertainties, since there is less independent evidence for drift-free N_2H_4 tank pressure transducers, drift must be taken into account. As before, the spec for pressure transducer drift was ±0.1% per year, up to ±2.0% total (presumably 3σ). This led to truly large uncertainties in N_2H_4 mass, particularly given tank pressure vs. consumed N_2H_4 mass follows a 1/x curve for a blowdown tank. Luckily, one further piece of information helped reduce the error contribution from potential pressure transducer drift—Figure 20 confirmed there was no relative drift between N_2H_4 tank and line pressure measurements. Therefore, even though pressure sensor drift was still possible, it was less likely to have occurred on both measurements. A Monte Carlo model was built to determine the joint probability of both sensors drifting in tandem. This decreased the a priori 3σ uncertainty to the equivalent of only 1.90σ for each transducer, so the pressure sensor drift error in Table 21 included a multiplicative reduction factor of 1.90/3.00. Despite this pencil sharpening, drift could still have caused a very large ±5.824 kg (3σ) N_2H_4 error.

The final term in Table 21 to be covered is tank pressure uncertainty not including transducer drift. As before, tank pressure uncertainty included 3σ spec uncertainties for linearity (±0.5%), hysteresis (±0.2%), repeatability (±0.1%), and compensated temperature range (±2.0%). Like for the BTA, the latter term was reduced via linear scaling, but this time by the final tank temperature uncertainty range (±5°C, 3σ). This reduced the ±2.0% (3σ) spec to only ±0.24% (3σ). These terms were then RSS'd with the assumed worst case ±1% A-to-D error, along with the ± ½ DN uncertainty of the tank pressure itself. Finally, this RSS'd pressure measurement error was converted to N_2H_4 mass error at hydrazine depletion, leading to the second largest term in Table 21, ±5.367 kg (3σ).

All sources of statistically unusable N_2H_4 were RSS'd to calculate the total of ±8.037 kg (3σ), or 6.1% of the launch load. This vastly exceeded the placeholder value of 1.0%, even before adding in deterministically unusable N_2H_4. With this piece included, Cassini's total unusable N_2H_4 mass at hypothetical tank depletion was a mammoth ±11.685 kg (3σ), or nearly 9% of the launch load of 132 kg. Perhaps unusable N_2H_4 for lengthy missions may be far higher than typically thought? Thankfully, Cassini's hydrazine margins ended up well above these uncertainties.

Table 21: Cassini Unusable Monopropellant Mass (Hold-Up) and Reserves Summary Table

Deterministically Unusable Monopropellant	**N_2H_4 [kg]**	**N_2H_4 [%]**
Trapped Liquid in Propellant Lines	0.608	0.461
Trapped Vapor in Propellant Tank @ Depletion	0.004	0.003
Trapped Due to Diaphragm Permeation, Diffusion, & Condensation	0.939	0.711
Trapped Due to Assumed 99.5% Tank Expulsion Efficiency	0.660	0.500
Unmodeled Usage due to STV Tank Temp Sensor Offset (reserve)	0.650	0.492
Unmodeled Usage due to Tank/Line Pressure Sensor Offset (reserve)	0.787	0.596
Total Deterministic	3.648	2.764
Statistically Unusable Monopropellant (3σ) in Reserves	**N_2H_4 [kg]**	**N_2H_4 [%]**
Hydrazine Loading Uncertainty	0.911	0.690
Tank Pressure Uncertainty @ Depletion (w/o Transducer Drift)	5.367	4.066
Tank Pressure Uncertainty @ Depletion due to Pressure Sensor Drift	5.824	4.412
Tank Temperature Uncertainty @ Depletion	0.531	0.402
Monopropellant System Helium Mass Uncertainty @ Depletion	0.869	0.658
Total Statistical (3σ, RSS) in Reserves	8.037	6.089
TOTAL UNUSABLE N_2H_4 (HOLD-UP) & RESERVES (3σ)	11.685	8.852

XIII. Conclusions

The Cassini mission was an ambitious voyage into the unknown, and it provided a chance to reveal the Saturnian system in detail for the first time. Over two decades, spacecraft performance was virtually flawless, including the complex bipropellant and monopropellant propulsion module subsystems. With the minor exception of early mission prime regulator leakage, slightly out-of-spec latch valve leakage, and potentially low check valve reseat pressures, the bipropellant CPMS functioned epically, well beyond its intended design life. Redundant systems, such as back-up MR-103H thrusters, proved their mettle and demonstrated their criticality for missions which may be extended beyond even the most optimistic timelines. Even though the ringed planet and its environs summon humanity's return with future missions, there may never be another spacecraft so consistently productive and trouble-free as Cassini.

During thirteen years in Saturn orbit, the spacecraft twice had the opportunity to image Saturn while the ringed planet eclipsed the sun (see Figure 66). Both times this "Kodak moment" happened, a single pale-blue pixel was readily visible within the ghostly lit background. Gazing upon our precious and precarious Earth from 900 million miles away reminded the author of the famous quote from T. S. Eliot's "Little Gidding":

"We shall not cease from exploration
And the end of all our exploring
Will be to arrive where we started
And know the place for the first time."

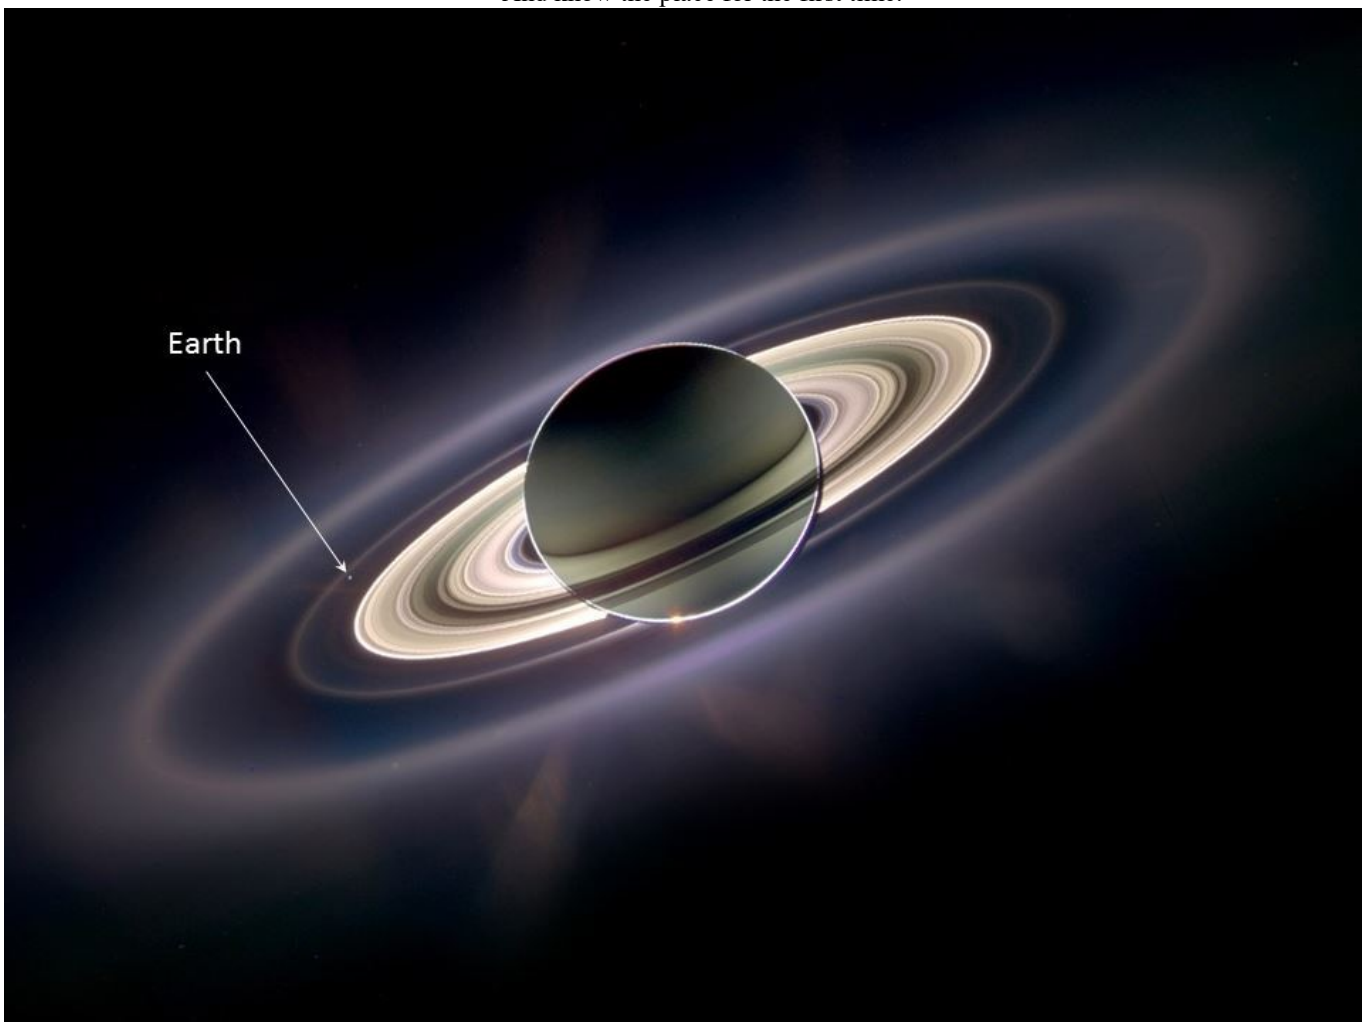

Figure 66: Cassini Image of Earth while Saturn Occulted the Sun

Acknowledgments

For a flagship mission as comprehensive and long-lived as Cassini, it is an impossible task to properly recognize the efforts of many thousands of engineers, scientists, technicians, managers, and support staff. Each and every member of the Cassini team made important contributions to its ultimate success and scientific triumphs at Saturn. In particular, though, the author would like to thank NASA's Space Science Directorate, the Executive Council at JPL, the stellar propulsion development team at JPL and Lockheed Martin Astronautics, the many propulsion component vendors, and the always supportive Propulsion Section at JPL. In addition, the author is most grateful for twenty glorious years working with excellent Cassini mission operations personnel, especially past and current Cassini Program Managers, the Spacecraft Operations Team Chief, fellow subsystems and systems engineers, the Cassini Outreach Team, the Cassini Mission Planning and Navigation teams, hundreds of science team members from around the globe, and the author's dedicated propulsion team. Guiding this incredible spacecraft for two decades will likely be the highlight of this author's entire career at JPL.

Dedication

The author wishes to dedicate this work to the memory of Mr. Richard T. Cowley, a wonderful mentor in propulsion mission operations at JPL. Dick was an invaluable Cassini propulsion colleague for decades and a dear personal friend. His many contributions to Cassini's success will not be forgotten.

References

[1] Barber, Todd J.; Krug, F. A.; and Froidevaux, B. M.: "Initial Galileo Propulsion System In-Flight Characterization," AIAA Paper 93-2117, June, 1993.

[2] Buffington, Brent; Strange, Nathan; and Smith, John: "Overview of the Cassini Extended Mission Trajectory," AIAA Paper 2008-6752, August, 2008.

[3] Barber, Todd J.; Krug, F. A.; and Renner, K. P.: "Final Galileo Propulsion System In-Flight Characterization," AIAA Paper 97-2946, July, 1997.

[4] Smith, John and Buffington, Brent: "Overview of the Cassini Solstice Mission Trajectory," AAS Paper 09-351, August 9-13, 2009.

[5] Buffington, Brent: "Proposed End-of-Mission for the Cassini Spacecraft: Inner D Ring Ballistic Saturn Impact," IAC Paper IAC-10-C1.9.2, September 27 – October 1, 2010.

[6] Barber, Todd J. and Cowley, R. T..: "Initial Cassini Propulsion System In-Flight Characterization," AIAA Paper 2002-4152, July 7-10, 2002.

[7] Leeds, Michael W.; Eberhardt, R. N.; and Berry, Robert L.: "Development of the Cassini Spacecraft Propulsion System," AIAA Paper 96-2864, July, 1996.

[8] Mizukami, M.; Barber, T. J.; Christodoulou, L. N.; Guernsey, C. S; and Haney, W. A.: "Cassini Spacecraft Reaction Control System Thrusters Flight Experience," JANNAF Paper (ITAR restricted), May, 2010.

[9] Goodson, Troy: "Evaluation of an Energy-Cutoff Algorithm for the Saturn Orbit Insertion Burn of the Cassini-Huygens Mission," AAS Paper 04-133, February 8-12, 2004.

[10] Brown, Todd S. : "Y-Biasing: A New Operational Method for Cassini to Control Thruster Usage while Managing Reaction Wheel Momentum," AIAA Paper 2010-7650, AIAA Guidance, Navigation, and Control Conference, 2-5 August 2010, Toronto, Ontario, Canada.

[11] Bates, David; Lee, Allan; Meakin, Peter; and Weitl, Raquel : "Fault Protection Design and Testing For the Cassini Spacecraft in a "Mixed" Thruster Configuration," AIAA Paper 2013-4632, AIAA Guidance, Navigation, and Control Conference, 19-22 August 2013, Boston, MA.

[12] Lee, Clifford C. and Lee, Allan Y: "Cassini Reaction Wheel Momentum Bias Optimization Tool," AIAA Paper 2005-6271, AIAA Guidance, Navigation, and Control Conference and Exhibit, 15-18 August 2005, San Francisco, CA.

[13] Sturm, Erick J.; Barber, Todd J.; and Roth, Duane: "Ensuring Cassini's end-of-mission propellant margins," IEEE 978-1-4799-5380-6, 2015 IEEE Aerospace Conference, 7-14 March, 2015, Big Sky, MT.

[14] Stephenson, R. Rhoads; Almaguer, Teofilo A., Jr.; Bernard, Douglas E.; Eisenman, David J.; Gindorf, Thomas E.; Guernsey, Carl S.; Lou, Michael C.; MacPherson, Duncan; Savino, Joseph L.; Slonski, John P, Jr.; Wood, Gordon E.; and Wright, Larry: "Mars Observer Loss of Signal: Special Review Board Final Report," JPL Publication 93-28, November, 1993.

[15] Kawakatsu, Yasuhiro; Campagnola, Stefano; Hirose, Chikako; and Ishii, Nobuaki: "An Orbit Plan Toward Akatsuki Venus Reencounter and Orbit Injection," AAS-12-205, Austin, TX, 8-12 January, 2012.

[16] Morgan, Paula S.: "Cassini Spacecraft's In-Flight Fault Protection Redesign for Unexpected Regulator Malfunction," AERO.2010.5446845, 2010 IEEE Aerospace Conference, Big Sky, MT, 2010, pp. 1-14.

